arXiv:1701.06386v1 [cs.CC] 23 Jan 2017

A Structured View on Weighted Counting with Relations to Counting, Quantum Computation and Applications

Cassio P. de Campos [*1], Georgios Stamoulis [†2], and Dennis Weyland [‡3]

[1]Queen's University Belfast, Northern Ireland, United Kingdom
[2]Department of Data Science & Knowledge Engineering, Maastricht University, The Netherlands
[3]Department of Economics and Management, University of Brescia, Italy

Abstract

Weighted counting problems are a natural generalization of counting problems where a weight is associated with every computational path of non-deterministic Turing machines and the goal is to compute the *sum of the weights of all paths* (instead of just computing the *number* of accepting paths). Many useful closure properties and plenty of applications make weighted counting problems interesting. The general definition of these problems captures even undecidable problems, but it turns out that obtaining an exponentially small additive approximation is just as hard as solving conventional counting problems. In many cases such an approximation is sufficient and working with weighted counting problems tends to be very convenient. We present a structured view on weighted counting by defining classes that depend on the range of the function that assigns weights to paths and by showing the relationships between these different classes. These classes constitute generalizations of the usual counting problems. Weighted counting allows us to easily cast a number of famous results of computational complexity in its terms, especially regarding counting and quantum computation. Moreover, these classes are flexible enough and capture the complexity of various problems in fields such as probabilistic graphical models and stochastic combinatorial optimization. Using the weighted counting terminology and our results, we are able to greatly simplify and answer some open questions in those fields.

1 Introduction

Counting problems play an important role in computer science and can be informally defined as finding the number of solutions (that is, computational paths of polynomial length that end in an accepting state) of a given combinatorial decision problem. Weighted counting problems are a natural generalization of counting problems: a weight is associated with every computational path and the goal is to compute the sum of the weights of all computational paths. Weighted counting has numerous applications in a wide variety of fields, such as quantum computation, stochastic combinatorial optimization, probabilistic graphical models, to name but a few. In many situations it is more natural and more convenient to use weighted counting problems instead of conventional counting problems. For instance, certain proofs can be simplified and stated in a more intuitive manner by using weighted counting. It also offers additional closure properties which do not hold

[*]E-mail: c.decampos@qub.ac.uk
[†]E-mail: georgios.stamoulis@maastrichtuniversity.nl
[‡]Email: dennisweyland@gmail.com

for the conventional counting classes. Because of that and because the computational complexity of weighted counting is closely related to that of conventional counting, the former might be preferred to the latter when studying some computational complexity questions.

In this work we give compact definitions for different classes of weighted counting problems, which differ only in the range of their weights. In particular, we study the complexity of weighted counting problems using natural, integer, rational and real numbers as weights together with some more restricted cases where weights take values on finite sets such as $\{0, 1\}$ and $\{-1, 1\}$. If one allows only positive integers as weights, it is easy to show that problems remain in #P, the class of (non-weighted) counting problems [57]. If negative values are allowed, even if they are integers, then it is not possible anymore to guarantee that the result is a non-negative value. Therefore, these problems are not in #P by definition. In order to study these cases, we adopt the terminology of #P-*equivalence* to indicate that a problem is #P-hard under 1-Turing reductions and can be affinely reduced to a problem in #P. This concept of affine reductions is closely related to weighted reductions [12], but it is not approximate preserving [26]. In short, it allows polynomial-time preprocessing of the input and an affine transformation of the output from the function. Essentially, the difference between #P-*equivalence* and previously used 1-Turing #P-*completeness* is the affine transformation for the membership in #P. Hence, problems which are not actually in #P because of technical details can still be #P-*equivalent*. Using this terminology, we show that problems using weights from the set $\{-1, 1\}$, and from arbitrary integers, are closely related to #P problems, and also that the decision variants of these weighted counting problems remain in PP (the class of problems related to the decision of whether the majority of the polynomial-time computational paths are accepting paths) [34]. While some of these results are considered to be known by the community, we are not aware of an explicit and precise treatment of all relevant cases.

The weighted counting problems just described are asking for computing results whose size is always bounded by a polynomial in the input size. In this case, the corresponding decision variants remain in PP. However, this is not necessarily the case for weighted counting problems when one allows weights to be taken from real or even rational numbers. In spite of that, we show that computing arbitrarily good additive approximations for these problems is a #P-equivalent problem. Regarding the corresponding decision problems, it is shown that the superclass of PP which allows rational weights (we will later denote it as $\mathrm{PP}_{\mathbb{Q}}$) can decide the halting problem, and hence strictly contains PP.

We also address the relations between complexity classes of weighted counting problems and those of quantum computation. We show that weighted counting can be used to represent the acceptance probability of a quantum Turing machine in a very natural way. Based on this observation, it is possible to derive complexity results regarding quantum computation using very simplified proofs. For instance, we are able to show that BQP $\subseteq$ AWPP [33] using a very simple and short argumentation (BQP is the class of decision problems that can be solved by a quantum Turing machine with bounded error probability [9] and AWPP is a quite restricted gap definable counting class and the best known classical upper bound for BQP [30]). Finally, we give a short and intuitive proof of the recently published result that quantum Turing machines with bounded error can be simulated in a time- and space-efficient way using randomized algorithms with unbounded error that have access to random access memory [59].

We conclude the paper with a discussion of the implications of our results for problems in other fields such as inferences in probabilistic graphical models and stochastic combinatorial optimization problems. Using these results and terminology, it is possible to simplify complexity proofs for those problems and to close some open questions, namely, upper bounds for the computational complexity of various computational tasks related to probabilistic graphical models and two-stage stochastic combinatorial optimization problems.

2 Counting Classes, Generalizations and Historical Developments

The starting point of our work is the definition of conventional counting problems and the corresponding class of decision problems. This definition is altered to allow for the summation of weights instead of counting accepting paths. Restrictions on the weights lead to several different classes of problems. We discuss basic properties of these classes and relate them to conventional counting problems. Additionally, we show that the closure properties for conventional counting problems can be extended to weighted counting problems. Although these properties make it convenient to work with weighted counting problems, the general definition leads to problems which might have an exponentially long output or even capture undecidable problems. We then show that from a complexity point of view approximating weighted counting problems up to an exponentially small additive error is basically equivalent to solving conventional counting problems. In many cases such an approximation is sufficient and therefore weighted counting is a convenient and powerful tool. We begin with an introduction to related concepts. Since counting classes play a central role in our paper, we devote part of this section to discuss them.

2.1 Counting Problems and Generalizations

Problems in the complexity class NP are related to the question of whether, for a given input, there exists at least one valid certificate of polynomial size (with respect to the input size) and which can also be checked in polynomial time (with respect to the input size). This complexity class has been generalized in 1979 [57, 58] in the following way: instead of deciding if there exists at least one valid certificate for a given input, we want to compute, for a given input, the number of valid certificates. Such problems are called counting problems and are subsumed in the complexity class #P. More formally, we can define counting problems as follows.

Definition 1 (Counting Problem)**.** *We are given a polynomial p and a polynomial-time predicate T (meaning a polynomial-time Turing machine whose output is either 0 or 1). The counting problem associated with p and T is to compute for $x \in \{0,1\}^\star$ the function*

$$f(x) = \sum_{u \in \{0,1\}^{p(|x|)}} T(x,u).$$

We also need definitions of related complexity classes (such as BPP, PP, BQP) while studying the relationships of weighted counting and other classes. Their definitions are given below.

Definition 2 (BPP)**.** *A language $L \subseteq \{0,1\}^\star$ is in the complexity class BPP if and only if there exists a polynomial p and a polynomial-time predicate M (meaning a polynomial-time Turing machine whose output is either 0 or 1) such that for every $x \in L$, the fraction of strings y of length $p(|x|)$ which satisfy $M(x,y) = 1$ is greater than or equal to $2/3$; for every $x \notin L$, the fraction of strings y of length $p(|x|)$ which satisfy $M(x,y) = 1$ is less than or equal to $1/3$.*[1]

Definition 3 (BQP)**.** *A language $L \subseteq \{0,1\}^\star$ is in the complexity class BQP if and only if there exists a polynomial-time uniform family of quantum circuits $\{Q_n : n \in \mathbb{N}\}$, such that:*

1. *For all $n \in \mathbb{N}$, Q_n takes n qubits as input and outputs 1 bit.*
2. *For all $x \in L$, $\Pr(Q_{|x|}(x) = 1) \geq \frac{2}{3}$.*
3. *For all $x \notin L$, $\Pr(Q_{|x|}(x) = 0) \geq \frac{2}{3}$.*

[1] The string y can be seen as the output of the random coin flips that the probabilistic Turing machine would have made.

This class is defined for a quantum computer and its natural corresponding class for an ordinary computer (or a Turing machine plus a source of randomness) is BPP. Just like P and BPP, BQP is low for itself, which means $\mathrm{BQP}^{\mathrm{BQP}} = \mathrm{BQP}$. In fact, BQP is low for PP, meaning that a PP machine (see the following definition) achieves no benefit from being able to solve BQP problems instantly, an indication of the possible difference in power between these related classes. The relation between BQP and NP is not known. Adding post-selection to BQP results in the complexity class PostBQP which is equal to PP [1].

The closest decision complexity class to #P is PP, standing for Probabilistic Polynomial time.

Definition 4 (PP)**.** *A language $L \subseteq \{0,1\}^\star$ is in the complexity class PP if and only if there exists a polynomial p and a polynomial-time predicate M (meaning a polynomial-time Turing machine whose output is either 0 or 1) such that:*

1. *For all $x \in L$, the fraction of strings y of length $p(|x|)$ which satisfy $M(x,y) = 1$ is strictly greater than $1/2$.*
2. *For all $x \notin L$, the fraction of strings y of length $p(|x|)$ which satisfy $M(x,y) = 1$ is less than or equal to $1/2$.*

As usual, the string y should be interpreted as a string of bits corresponding to random choices. The terms "less than or equal" can be changed to "less than" and the threshold $1/2$ can be replaced by any fixed rational number in $]0,1[$, without changing the class. PP is a very powerful class: it contains BPP, BQP and NP. In the following we give a new and extremely short intuitive proof that PP contains BQP (and also the related class AWPP, see Section 4) based entirely on the concept of weighted counting. Moreover, as we will discuss shortly, a polynomial-time Turing machine with access to a PP oracle can solve the entire polynomial hierarchy. On the other hand, PP is contained in PSPACE (polynomial-space bounded Turing Machines; for instance, polynomial space can solve the *majority satisfiability* problem simply by going through each possible assignment). Another concept from complexity theory that we need is the notion of *Functional Problems*.

Definition 5 (FP)**.** *A binary relation $P(x,y)$ is in the complexity class FP if and only if for every $x \in \{0,1\}^\star$ there is a deterministic polynomial-time (in $|x|$) algorithm that finds some $y \in \{0,1\}^\star$ such that $P(x,y)$ holds (or tells that such y does not exist).*

The difference between FP and P is that problems in P have one-bit yes/no answers, while problems in FP can have any output that can be computed in polynomial time. Just as P and FP are closely related, NP is closely related to FNP. In a straightforward way, the class of relations $\mathrm{FP}^{\#\mathrm{P}}$ contains all relations that can be computed in polynomial time with access to a #P oracle. We use the terminology $\mathrm{FP}^{\#\mathrm{P}[m]}$ to limit the overall number of calls to the oracle to m. We note that $\mathrm{FP}^{\#\mathrm{P}}$ is a very powerful class, as it was shown in [56]: every function in #PH is Turing-1 reducible in polynomial time to some function in #P, thus $\#\mathrm{PH} \subseteq \mathrm{FP}^{\#\mathrm{P}[1]}$ (where #PH is the class of functions that count the number of accepting paths of polynomial time-bounded nondeterministic oracle Turing machines with oracle sets from the polynomial hierarchy PH).

2.2 Related Work and Historical Developments

Since its introduction, the class #P has been very successful in characterizing the difficulty of counting problems: for this, central is the role of #P-completeness [57, 58] introduced to capture the computational complexity of the *permanent*. Surprisingly, not only all NP-complete problems have their counting version #P-complete (this is true under *parsimonious* reductions) but also many "easy" problems such as *perfect matching* are #P-complete in their counting versions. These problems with easy decision version (called "hard to count-easy to decide" in [23]) cannot be #P-complete under *parsimonious* reductions, but are #P-complete under Cook reductions. See [46] for some closely related classes such as TotP and #PE (standing for #P-easy) and their corresponding

structural properties and characterizations and also [38, 52] for work on some other closely related subclasses of #P that contain functions with easy decision variants.

In the classical "textbook" proof that computing the permanent of a matrix $M \in \mathbb{Z}^{n\times n}$ is #P-complete [8], first it shown a reduction to the computation of the permanent of another matrix $M' \in \mathbb{Z}_{\geq 0}^{n\times n}$ and then to the computation of the permanent of a matrix M'' with 0/1 entries. In other words, a counting problem over arbitrary integers is reduced to a classic 0/1 counting problem. This however requires pre and postprocessing, so it has been argued that #P might not be the correct class when we are interested in counting. Indeed #P has some disadvantages: the class #P is not closed under many operations. Namely, #P is not (or is not known to be) closed under subtractions or binomial coefficients. A direct implication of the first case is that computing the permanent of a matrix with arbitrary integer entries is not #P by definition. This was a main motivation behind the work of Fenner, Fortnow and Kurtz [30] that introduced and defined the complexity class GapP (intuitively, a class of functions that introduces "gaps" between the number of accepting and rejecting computations of NP machines) as an alternative to #P (see also [31] for a more systematic study of gap definability). The authors claim that this class constitutes a natural alternative for #P. Indeed, many computations (such as the permanent over arbitrary integers) which are outside #P fall now inside GapP. Moreover, GapP is shown to be closed under non-trivial operations such as subtractions, binomial coefficients, etc (see also [6]). Arguably, the success of GapP comes from the work of Beigel, Reingold and Spielman to show that PP is closed under the operation of intersection [7] and also due to the fact that the definition of GapP helps to simplify very important complexity theoretic results: Toda's famous theorem that the entire polynomial hierarchy is contained in P^{PP} [54] can be cast in terms of GapP [55]. The very important result that BQP $\subseteq$ PP by Adleman, DeMarrais and Huang [3] can also be simplified and improved with GapP: in [33], among others, it has been shown, using GapP definable functions, that $\mathrm{PP}^{\mathrm{BQP}} = \mathrm{PP}$. Moreover, it has been shown that BQP $\subseteq$ AWPP, which is the best current upper bound for BQP. Regarding the simulation of a BQP machine by a randomized, unbounded error machine, the most efficient simulation was given in a very recent article by van Malkebeek and Watson [59]. In spite of that, it seems that the class GapP cannot capture the complexity of *weighted counting*, where every computational path has a weight and we are interested in computing the sum of weights of all computational paths, which is precisely the focus of this work.

In another line of research, the complexity of counting *constraint satisfaction problems* (CSPs) has been investigated. We recall that a CSP can be formulated as follows: Let D be an arbitrary finite set called the domain set. Let $\mathcal{R} = \{R_1, R_2, \ldots, R_k\}$ be an arbitrary finite set of relations on D. Each R_i has an arity $n_i \in \mathbb{N}$. As input we have variables $x_1, \ldots, x_n$ over D, and a collection of constraints $R \in \mathcal{R}$, each applied to a subset of variables. The decision query is whether there is an assignment that satisfies all the constraints and the corresponding counting version asks to compute the number of satisfying assignments. First, identify each $R \in \mathcal{R}$ with a binary valued function $\phi(R)$. Then a counting CSP is to evaluate the following so-called partition function on an input instance I:

$$P(I) = \sum_{x_i \in D} \prod_{f \in I} f(x_{i_1}, \ldots, x_{i_r}),$$

where f has arity r and is applied to variables $x_{i_1}, \ldots, x_{i_r}$. If f is 0/1 valued, then this counts the number of solutions. If ϕ can take arbitrary values, then we have a *weighted* CSP problem (see for example [12] and [24]).

Feder and Vardi [28] conjectured that a CSP over constraints $\mathcal{R}$ is either in P or NP-hard, and very recent work may have settled this question [48]. In a very recent breakthrough, Cai and Chen [15, 16] proved a dichotomy theorem for the counting versions of CSPs over complex weights: Given any finite D and any finite set of constraint functions $\mathcal{F}$, where $f_i : D^{n_i} \to \mathbb{C}$, the problem of computing the partition function $P(I)$ is either in P or #P-hard. On the negative side, the criteria are not known to be computable, but still the dichotomy exists. Another similar dichotomy conjecture, but yet to be proved, exists for #CSP, namely that any #CSP problem is

either in FP or #P-complete. See [13] for some partial results regarding this conjecture and see also [24, 17, 12, 14, 63, 64] for other results in the are of counting versions of CSPs.

Another related result comes from Bläser and Curticapean [11] in which the #W[1]-hardness of the weighted counting of all k-matchings in a bipartite graph was given. There, the weight of a matching is simply the product of the weights of the edges that constitute this matching.

In yet another line of research and more relevant to our study, a weighted counting class $\#\mathrm{P}_{Q'}$ has been defined by Goldberg and Jerrum [35]. This is the class of functions $f : \{0,1\}^* \to \mathbb{Q}$ that can be written as the division of a #P function over an FP function. This class was used to classify the complexity of approximating the value of Tutte polynomials of a graph that take as argument arbitrary rational numbers. Recall that the Tutte polynomial of a given graph G is a two-variable polynomial $T_G(x, y)$. Usually, the arguments x, y are integers (not necessarily positive), but the authors were interested in the most general case where x, y are arbitrary rational numbers. Tutte polynomials can encode interesting graph theoretic properties. For instance, $T_G(1, 1)$ counts the number of spanning trees in G, $T_G(2, 1)$ counts the number of forests in G, and so on. Tutte polynomials are also very closely connected to chromatic polynomials, flow polynomials, etc. This class $\#\mathrm{P}_{Q'}$ is a strict subclass of the class $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ defined later on. We shall mention that the complexity of exactly evaluating Tutte polynomials is #P-hard [39] except for some threshold cases. The authors were interested in some dichotomy results and they significantly widened the cases where there exist a fully polynomial time randomized approximation scheme (FPRAS) for $T_G(x, y)$, and showed that some other particular cases do not attain FPRAS (modulo RP $\neq$ NP).

Finally, it is important to mention the result of Yamakami [62] that is probably the closest to the settings in this paper. There, the author studied a class of quantum functions, namely $\#\mathrm{QP}_{\mathrm{K}}$, which are defined as the set of functions computing the acceptance probability of some polynomial-time quantum Turing machine, the amplitudes of which are drawn from a set K. We follow a similar notation here. He also defined the corresponding gap definable function $\mathrm{GapQP}_{\mathrm{K}} = \#\mathrm{QP}_{\mathrm{K}} - \#\mathrm{QP}_{\mathrm{K}}$. He named these functions as quantum probability and quantum probability gap functions, respectively. Among other very interesting results, Yamakami proved the following nice characterization of PP in terms of some generalized quantum classes:

Theorem 1 (Theorem 7.1 in [62]). *Let $A \subseteq \{0,1\}^*$ and let $\mathbb{A}$ be the set of complex algebraic numbers. Then, the following statements are all equivalent:*

1. *$A \in PP$,*
2. *$\exists f, g \in \#\mathrm{QP}_{\mathbb{A}}$ such that $\forall x \in \{0,1\}^*$, $x \in A \Leftrightarrow f(x) > g(x)$,*
3. *$\exists f, g \in \mathrm{GapQP}_{\mathbb{A}}$ such that $\forall x \in \{0,1\}^*$, $x \in A \Leftrightarrow f(x) > g(x)$.*

The usefulness of this result is even more evident when Yamakami, in the same paper, considers the quantum analog of PP, namely $\mathrm{PQP}_{\mathrm{K}}$ for some amplitude set K (where the usual polynomial-time Turing machine of the standard definition of PP is replaced by a certain quantum Turing machine). Using the above theorem, Yamakami proved that $\mathrm{PQP}_{\mathbb{A}} = \mathrm{PP}$. Later in this paper, we will show that PP over arbitrary (approximable) rational numbers contains even *undecidable* problems and thus cannot be equal to PP. However, we do not know if the statement is true if we consider a more restricted definition.

2.3 Weighted Counting Problems

As a straightforward generalization of conventional counting problems, weighted counting problems can be formally defined in the following way.

Definition 6 (Weighted Counting Problem). *We are given a polynomial p and a function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{R}$ that can be approximated by a polynomial-time (in the size of the first two arguments and the third argument) computable function $v : \{0,1\}^\star \times \{0,1\}^\star \times \mathbb{N} \to \mathbb{Z}$, such that*

$|w(x,u) - v(x,u,b)/2^b| \leq 2^{-b}$ *for all* $x \in \{0,1\}^\star, u \in \{0,1\}^\star, b \in \mathbb{N}$. *The weighted counting problem associated with* p *and* w *is to compute for* $x \in \{0,1\}^\star$ *the function*

$$f(x) = \sum_{u \in \{0,1\}^{p(|x|)}} w(x,u).$$

Here, x is the input to the weighted counting problem and w represents the $2^{p(|x|)}$ many weights for a given input x. With the restriction to w, imposed by the approximation property, we limit the range of w to efficiently computable numbers as defined in [44]. A similar notion in the context of quantum computation is used in [3, 59]. For a given rational threshold value we can then define the corresponding decision problem in the following way.

Definition 7 (Weighted Counting Problem, Decision Variant)**.** *We are given a weighted counting problem defined by a polynomial* p *and a function* $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{R}$ *as well as a threshold value* $t \in \mathbb{Q}$. *The corresponding decision problem is to decide for* $x \in \{0,1\}^\star$ *whether* $f(x) \geq t$ *or not.*

Without loss of generality, we could have used an integer threshold value in the above definition, since the denominator of the threshold value can always be multiplied into the weight function. Nevertheless, the definition given above might be more convenient for applications of weighted counting, while assuming an integer threshold value could be used to simplify proofs.

As mentioned earlier, weighted counting problems may have different characteristics depending on the set from which weights are taken. For any given set $S \subseteq \mathbb{R}$, we define the class $\#\mathrm{P}_S$ to consist of all functions $f : \{0,1\}^\star \to \mathbb{R}$ corresponding to weighted counting problems where the range of w is restricted to S. The class of the corresponding decision problems is then denoted by PP_S. By using this notation, we can define the classes $\#\mathrm{P}_{\{0,1\}}$, $\#\mathrm{P}_{\{-1,1\}}$, $\#\mathrm{P}_{\{-1,0,1\}}$, $\#\mathrm{P}_{\mathbb{N}}$, $\#\mathrm{P}_{\mathbb{Z}}$, $\#\mathrm{P}_{\mathbb{Q}}$ and $\#\mathrm{P}_{\mathbb{R}}$, as well as the corresponding classes of decision problems $\mathrm{PP}_{\{0,1\}}$, $\mathrm{PP}_{\{-1,1\}}$, $\mathrm{PP}_{\{-1,0,1\}}$, $\mathrm{PP}_{\mathbb{N}}$, $\mathrm{PP}_{\mathbb{Z}}$, $\mathrm{PP}_{\mathbb{Q}}$ and $\mathrm{PP}_{\mathbb{R}}$. These classes contain interesting problems, since for instance the *permanent* of matrices with entries that are not natural numbers already falls in such different classes depending on the range of the entries. Some inclusions among these classes are trivial, since $S_1 \subseteq S_2 \Rightarrow (\#\mathrm{P}_{S_1} \subseteq \#\mathrm{P}_{S_2}) \wedge (\mathrm{PP}_{S_1} \subseteq \mathrm{PP}_{S_2})$.

It is easy to see that $\#\mathrm{P}_{\{0,1\}}$ is equal to #P [57], the class of (non-weighted) counting problems. The same equality holds for the corresponding classes of decision problems $\mathrm{PP}_{\{0,1\}}$ and PP. Additionally, $\#\mathrm{P}_{\{-1,0,1\}}$ is equal to GapP [30], the closure of #P under subtraction, while $\#\mathrm{P}_{\{-1,1\}}$ equals GapP under *normal* Turing machines (the computational tree is completely binary), and therefore $\mathrm{PP}_{\{-1,1\}}$ and $\mathrm{PP}_{\{-1,0,1\}}$ are equal to PP, since #P and GapP constitute the same class of decision problems.

The classes $\#\mathrm{P}_{\mathbb{R}}$ and $\mathrm{PP}_{\mathbb{R}}$ are extremely powerful since the output of $\#\mathrm{P}_{\mathbb{R}}$ can be infinitely long. For instance, $\mathrm{PP}_{\mathbb{R}}$ contains the *halting* problem. The following result confirms that not only $\#\mathrm{P}_{\mathbb{R}}$ and $\mathrm{PP}_{\mathbb{R}}$ but also $\#\mathrm{P}_{\mathbb{Q}}$ and $\mathrm{PP}_{\mathbb{Q}}$ are extremely powerful classes, since they also contain the *halting* problem. Let $\langle a, b \rangle$ represent the concatenation of the bitstring representations of a and b.

Theorem 2. *$\#P_{\mathbb{Q}}$ contains unsolvable problems and $PP_{\mathbb{Q}}$ can decide the halting problem.*

Proof. Define a weight function w that, on input $x = \langle M, y \rangle$, is as follows: $w(x,0) = 0$ if the Turing machine represented by M does not terminate on input y, while $w(x,0) = \frac{1}{2^t}$ if the Turing machine represented by M terminates after t steps on input y. We also define $w(x,u) = 0$ for every $u \neq 0$. As required by Definition 6, this weight function w can be approximated by a polynomial-time computation that on input $(x,0,b)$ simulates M on input y and if the simulation does not terminate after $b + p(|x|)$ steps, then it outputs 0, while if the simulation terminates at step $t \leq b + p(|x|)$, then it outputs $\frac{1}{2^t}$. The additive error of this approximation is at most 2^{-b}, since there are less than or equal to $2^{p(|x|)}$ paths with an approximate value, each of which with an error of at most $2^{-b-p(|x|)}$. Therefore, we have constructed a problem in $\#\mathrm{P}_{\mathbb{Q}}$ such that M terminates on input y if and only if $f(x) > 0$ and thus if we were able to decide the corresponding counting problem, then we would be able to decide the *halting* problem. □

Corollary 1. *PP is strictly contained in $PP_{\mathbb{Q}}$.*

Theorem 2 suggests that our definition gives too much power to $\mathrm{PP}_{\mathbb{Q}}$ and $\#\mathrm{P}_{\mathbb{Q}}$ and they are possibly equal to $\mathrm{PP}_{\mathbb{R}}$ and $\#\mathrm{P}_{\mathbb{R}}$, respectively. Since we are interested in understanding the complexity of practical problems, as we will discuss in Section 5, we decide to work with a restricted version of $\mathrm{PP}_{\mathbb{Q}}$ and $\#\mathrm{P}_{\mathbb{Q}}$ from now on.

Definition 8. *$\#P_{\mathbb{Q}}^{(\mathrm{pt})}$ is the subset of $\#P_{\mathbb{Q}}$ for which the weighted counting problems have their function w (as in Definition 6) computable in polynomial time (in the size of its input). $PP_{\mathbb{Q}}^{(\mathrm{pt})}$ is the associated decision version.*

In the following we further investigate the relations among the other classes of weighted counting problems and their corresponding decision versions. Before proceeding, we must define the terminology of #P-equivalence that will be used throughout the paper. This definition was necessary because #P is not closed with respect to polynomial-time computations. This leads to several inconveniences. We make use of an adapted version of weighted reduction [12], as suggested in [20], but we allow for affine postprocessing of the output from the desired function. For other definitions of metric reductions, see for example [27].

Definition 9. *A function $f : \{0,1\}^\star \to \mathbb{R}$ affinely reduces to $g : \{0,1\}^\star \to \mathbb{R}$ if there are polynomial-time computable functions h_1, h_2, h_3 such that $\forall x \in \{0,1\}^\star : f(x) = h_1(x) \cdot g(h_2(x)) + h_3(x)$.*

Definition 10. *A problem A is #P-hard under 1-Turing metric reductions if $\#P \subseteq FP^{A[1]}$. Problem A is said #P-complete if A is #P-hard and $A \in \#P$.*

Note that we use a restricted version of hardness when compared to the definition in Valiant's seminal work [58]. There, an unlimited number of oracle calls are allowed, that is, problem y is #P-hard if $\#\mathrm{P} \subseteq \mathrm{FP}^y$. This definition with unlimited calls allows PP-complete problems such as *majority satisfiability* to be #P-complete. This is somewhat undesired, since the classes PP and #P are not known to have similar power. For instance, $\mathrm{PP} = \mathrm{P}^{\mathrm{PP}[log]}$ [7, 32], while the corresponding question for #P is open, to the best of our knowledge. In particular, it is unknown whether $\mathrm{FP}^{\#\mathrm{P}[1]}$ is strictly contained in $\mathrm{FP}^{\#\mathrm{P}[2]}$ [36, 49]. As discussed earlier, $\#\mathrm{PH} \subseteq \mathrm{FP}^{\#\mathrm{P}[1]}$, suggesting that limited calls to #P are more powerful than limited calls to PP, while unlimited calls give us no distinction between them and $\mathrm{FP}^{\mathrm{PP}} = \mathrm{FP}^{\#\mathrm{P}}$. We note that this distinction may have practical implications. For instance, the recently very active topic of *probabilistic logics* has seen some important problems falling into classes such as $\mathrm{PP}^{\mathrm{NP}}, \ldots, \mathrm{PP}^{\mathrm{PH}}$ [18, 19]. All those classes are confined between the limited and unlimited PP calls: $\mathrm{PP} = \mathrm{P}^{\mathrm{PP}[log]} \subseteq \mathrm{PP}^{\mathrm{NP}} \subseteq \mathrm{PP}^{\mathrm{PH}} \subseteq \mathrm{P}^{\mathrm{PP}}$ [54].

Definition 11. *A problem A is #P-equivalent if it is #P-hard under 1-Turing metric reductions and can be affinely reduced to a problem in #P.*

The concept of #P-equivalence for problems relaxes the concept of #P-completeness under 1-Turing reductions by introducing the affine reduction instead of actual membership. Arguably, such concept gives a more useful relation between problems that are said to be #P-equivalent when compared to #P-completeness. Any #P-complete under 1-Turing reduction problem is also #P-equivalent (the inverse does not necessarily hold; for instance, *permanent* of integer matrices is #P-equivalent but is not #P-complete, since it does not belong to #P). Actually, the derivations we have in this paper would work even if we restricted ourselves to affine reductions for showing hardness, but 1-Turing reductions might be a reasonable compromise.

Definition 12. *A class C is #P-equivalent if any problem in C can be reduced to a problem in #P via an affine reduction and vice-versa, that is, any problem in #P can be reduced to a problem in C via an affine reduction.*

The notion of equivalence is commutative, and could be applied to any other functional complexity class. It gives a means to represent the strong relationship between some classes, as we

will see later on. It is worth noting that the decision variants of problems in #P-equivalent classes are in PP.

We first show equality of the classes $\#\mathrm{P}_{\{0,1\}}$ and $\#\mathrm{P}_{\mathbb{N}}$, as well as equality of the classes $\#\mathrm{P}_{\{-1,0,1\}}$ and $\#\mathrm{P}_{\mathbb{Z}}$. This implies equality of the corresponding classes of decision problems, namely of $\mathrm{PP}_{\{0,1\}}$ and $\mathrm{PP}_{\mathbb{N}}$, and of $\mathrm{PP}_{\{-1,1\}}$, $\mathrm{PP}_{\{-1,0,1\}}$ and $\mathrm{PP}_{\mathbb{Z}}$. These results are widely understood as "known" by the community, but to the best of our knowledge they have never been explicitly stated. Since the range of the functions in $\#\mathrm{P}_{\{-1,1\}}$ is not limited to non-negative integers, as it is the case for the functions in $\#\mathrm{P}_{\{0,1\}}$, these two classes cannot be equal. Nevertheless, it is possible to show that all these classes are #P-equivalent.

We later focus on weighted counting problems in $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$, $\#\mathrm{P}_{\mathbb{Q}}$ and $\#\mathrm{P}_{\mathbb{R}}$ as well as on their corresponding decision variants. $\#\mathrm{P}_{\mathbb{Q}}$ and $\#\mathrm{P}_{\mathbb{R}}$ capture undecidable problems, as we have proven. It is immediate to show that $\#\mathrm{P}_{\mathbb{Z}} \subseteq \#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ and $\mathrm{PP}_{\mathbb{Z}} \subseteq \mathrm{PP}_{\mathbb{Q}}^{(\mathrm{pt})}$. Additionally, the size of the output of weighted counting problems belonging to $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ is not necessarily polynomially bounded by the input size. Therefore, the inclusion of $\#\mathrm{P}_{\mathbb{Z}}$ in $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ is strict, and it is generally not possible to give polynomial reductions from problems in $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ to any of the more restricted classes. As we will see in Section 2.5, it is still possible to solve problems in $\#\mathrm{P}_{\mathbb{R}}$ using a #P-equivalent problem such that we lose only a small additive approximation error. This result then implies that $\mathrm{PP}_{\mathbb{Q}}^{(\mathrm{pt})}$, $\mathrm{PP}_{\mathbb{Q}}$ and $\mathrm{PP}_{\mathbb{R}}$ are actually PP as long as we focus on problems whose output size is polynomially bounded in the input size.

We make use of the following property to show equality between $\#\mathrm{P}_{\{0,1\}}$ and $\#\mathrm{P}_{\mathbb{N}}$. Given the weight function w and its approximation v (Definition 6), we can assure that the integer part of the weights can always be encoded using polynomially many bits. We then add this polynomial to the given polynomial p and construct a new weight function using only weights of 0 and 1, such that the overall sum does not change. We formalize these ideas in the proof of the following theorem.

Theorem 3. *$\#P_{\{0,1\}} = \#P_{\mathbb{N}}$.*

Proof. Since problems in $\#\mathrm{P}_{\{0,1\}}$ are by definition in $\#\mathrm{P}_{\mathbb{N}}$, we only have to show the other inclusion. For a weighted counting problem in $\#\mathrm{P}_{\mathbb{N}}$ defined by a polynomial p and a weight function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{N}$, we construct the following weighted counting problem in $\#\mathrm{P}_{\{0,1\}}$. Take the polynomial $p' = p + q$, where q is a polynomial bounding the number of bits required to encode the integer part of the weights of our original problem. Define $w' : \{0,1\}^\star \times \{0,1\}^\star \to \{0,1\}$ by

$$w'(x,u) = \begin{cases} 1 & \text{if } \#u_2 < w(x,u_1) \\ 0 & \text{else,} \end{cases}$$

where u_1 are the first $p(|x|)$ bits of u and $\#u_2$ is the number encoded by the last $q(|x|)$ bits of u. Since the two functions defined by the original weighted counting problem and the newly constructed weighted counting problem are identical, we conclude the proof. □

Theorem 4. *$\#P_{\{-1,0,1\}} = \#P_{\mathbb{Z}}$.*

Proof. It follows from the same arguments as in the proof of the previous theorem, this time applied to positive and negative weights. □

Observe that $\#\mathrm{P}_{\{-1,1\}}$ corresponds to GapP functions on normal Turing Machines (the computational tree is completely binary) whereas $\#\mathrm{P}_{\{-1,0,1\}} = \#\mathrm{P}_{\mathbb{Z}}$ corresponds to GapP functions. Since GapP strictly contains normal GapP functions, we immediately have $\#\mathrm{P}_{\{-1,1\}} \subset \#\mathrm{P}_{\{-1,0,1\}}$.

The classes $\#\mathrm{P} = \#\mathrm{P}_{\{0,1\}} = \#\mathrm{P}_{\mathbb{N}}$, $\#\mathrm{P}_{\{-1,1\}}$ and $\#\mathrm{P}_{\{-1,0,1\}} = \#\mathrm{P}_{\mathbb{Z}}$ cannot be equal, since their ranges of output are different. Nevertheless, there are affine reductions from problems of each of these classes to the other classes.

Theorem 5. *$\#P_{\{-1,1\}}$ and $\#P_{\{-1,0,1\}}$ (which equals to $\#P_{\mathbb{Z}}$) are #P-equivalent.*

Proof. Since $\#\mathrm{P}_{\{-1,1\}} \subset \#\mathrm{P}_{\{-1,0,1\}}$, it is enough to show reductions from #P to $\#\mathrm{P}_{\{-1,1\}}$ and from $\#\mathrm{P}_{\{-1,0,1\}}$ to #P. Any problem in $\#\mathrm{P}_{\{0,1\}}$ (which equals #P) can be affinely reduced to $\#\mathrm{P}_{\{-1,1\}}$ by multiplying all its weights by 2 and then subtracting 1. After solving the problem in $\#\mathrm{P}_{\{-1,1\}}$, the result is retrieved back using the inverse transformation, that is, dividing the output by 2 and summing $2^{p(|x|)-1}$. Any problem in $\#\mathrm{P}_{\{-1,0,1\}}$ can be affinely reduced to $\#\mathrm{P}_{\mathbb{N}}$ (which equals #P). For that we can simply add a value of 1 to every weight of the given $\#\mathrm{P}_{\{-1,0,1\}}$ problem, obtaining a problem in $\#\mathrm{P}_{\mathbb{N}}$ with weights in $\{0,1,2\}$. We can then retrieve the result of the original problem in $\#\mathrm{P}_{\{-1,0,1\}}$ by solving this new problem and then subtracting $2^{p(|x|)}$ from the resulting value. □

This relation of equivalence clarifies that, at least from a computational viewpoint, problems that are #P-equivalent are not harder to solve than those in #P, as a single evaluation of a #P function plus an affine transformation gives us the desired result. This fact might not be so surprising, but it seems that the literature still has not taken it for granted. For instance, Chapter 17.3 of [4] suggests to the reader an approach to solve *permanent* of integer matrices using unlimited calls to a #P oracle (more precisely, a #SAT one). It also suggests two calls to solve matrices with entries in $\{-1,0,1\}$. With the weighted counting framework, it becomes clear that a single call is enough, even if entries would be drawn from rational numbers (as we will see later on – the key is that the output of permanent is always polynomially bounded in the input size), and that the problem is in $\mathrm{FP}^{\#\mathrm{P}[1]}$, perhaps also in some lower class (any #P-equivalent problem is trivially in $\mathrm{FP}^{\#\mathrm{P}[1]}$).

In the remainder of this section, we show a proof that for weighted counting problems in $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ the size of the output, even encoded as a fraction, is not necessarily bounded by a polynomial in the input size. This proves that $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ cannot be #P-equivalent. We focus on the following simple problem. According to Definition 6, take the polynomial p to be the identity function and take the weight function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{Q}$ such that

$$w(x,u) = \begin{cases} 1/(\#u+1), & \text{if } \#u+1 \in \mathbb{P} \text{ and } \#u < x \\ 0, & \text{otherwise.} \end{cases}$$

Here $\#u$ is the number represented by the bitstring u and $\mathbb{P}$ is the set of prime numbers. Note that this function fulfills the approximation properties of Definition 6. For a given input $x \in \mathbb{N}$ of size $\lceil \log x \rceil$, the task is to compute the value

$$f(x) = \sum_{\substack{p \in \mathbb{P} \\ p \le x}} 1/p = \left(\sum_{\substack{p \in \mathbb{P} \\ p \le x}} \prod_{\substack{q \in \mathbb{P} \\ q \le x, q \ne p}} q \right) / \prod_{\substack{p \in \mathbb{P} \\ p \le x}} p.$$

Note that this fraction cannot be simplified. The only numbers that divide the denominator are prime numbers of value at most x. Each of these prime numbers divides all parts of the sum except one and therefore it does not divide the whole sum.

To show that the result cannot be represented efficiently with respect to the input size, we show that it is not possible to represent the denominator efficiently with respect to the input size. For this purpose we present a lower bound for the product of all prime numbers between x/e and x for sufficiently large x. Using this lower bound we can then bound the whole product appearing in the denominator of the output from below.

Lemma 1. *The number of prime numbers between x/e and x is bounded away from below by $x/(3 \ln x)$ for sufficiently large x.*

Proof. Let $\pi(x)$ denote the prime-counting function. Due to [50], we have for $x \ge 17$ that

$$\pi(x) > \frac{x}{\ln x} \quad \text{and} \quad \pi(x/e) < 1.25506 \frac{x/e}{\ln(x/e)}.$$

For the number of prime numbers between x/e and x we then have

$$\pi(x) - \pi(x/e) > \frac{x}{\ln x} - 1.25506\frac{x/e}{\ln(x/e)} > \frac{x}{\ln x} - 1.25506\frac{x/e}{3/4 \cdot \ln x} > \frac{1}{3}\frac{x}{\ln x}.$$

Hence this claim holds for sufficiently large x. □

This means that the product of all the prime numbers between x/e and x is bounded from below by $(x/e)^{x/(3\ln x)}$. Unfortunately, this value is doubly exponential in the input size of $\lceil \log x \rceil$ and a representation of this value would require exponentially many bits in the input size. Note that the result computed by a conventional counting problem or a counting problem using integer weights can always be represented using only polynomially many bits in the input size. Hence, a polynomial reduction between these classes cannot exist in general. This result is summarized in the following theorem.

Theorem 6. *$\#P_{\mathbb{Q}}^{(\mathrm{pt})}$ is **not** $\#P$-equivalent and is* **not** *a subset of $FP^{\#P}$.*

Proof. Any problem in a #P-equivalent class, or in $\mathrm{FP}^{\#\mathrm{P}}$, has output size bounded by a polynomial of the input size. $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ contains problems which do not have bounded output size, so the result follows. □

2.4 Closure properties

Like conventional counting and GapP, weighted counting comes with many useful closure properties. By definition, weighted counting is closed under addition and multiplication with a constant. Additionally, it is closed under uniform sums of exponentially many terms and uniform products of polynomially many terms. Moreover, weighted counting is closed under finite sums and products. These properties basically follow from the definition of weighted counting and the proofs are analogous to those for conventional counting and GapP [6]. We summarize these properties in the following proposition.

Proposition 1. *Let $f_1, \ldots, f_k$, $k \in \mathbb{N}$, be functions computed by weighted counting problems. Then the following functions can be computed by a weighted counting problem.*

(a) $g(x) = f_1(x) + c$, *where c is a constant.*

(b) $g(x) = c \cdot f_1(x)$, *where c is a constant (and therefore in particular $-f_1(x)$).*

(c) $g(x) = \sum_{i=1}^{k} f_i(x)$.

(d) $g(x) = \prod_{i=1}^{k} f_i(x)$.

(e) $g(x) = \sum_{y \in \{0,1\}^{p(|x|)}} f_1(\langle x, y \rangle)$, *where p is a polynomial.*

(f) $g(x) = \prod_{y \leq p(|x|)} f_1(\langle x, y \rangle)$, *where p is a polynomial.*

Proof. These properties basically follow from the definition of weighted counting and the proofs are analogous to those for conventional counting and GapP [6]. We include here a discussion for completeness.

In order to prove that $g(x)$ can be computed by some weighted counting function, we need to specify the functions w and v (serving as the approximation of w) in accordance to Definition 6. Most of the cases are a very easy application of Definition 6. For example, for case (a) the function w' for g is the same as the function w of f_1 plus the term c. Since f_1 is a function computed by a weighted counting problem, there exist an approximation function v such that $|w(x,u) - v(x,u,b)/2^b| \leq 2^{-b}$ for all $x \in \{0,1\}^\star, u \in \{0,1\}^\star, b \in \mathbb{N}$. Define the approximation function v' for g to be $v' := v + c \cdot 2^b$ for $b \in \mathbb{N}$. Observe that v' is easily computable in polynomial time (e.g. repeated squaring to compute the term 2^b). Then we see that

$$|w'(x,u) - v'(x,u,b)/2^b| = |w(x,u) + c - v(x,u,b)/2^b - (c \cdot 2^b/2^b)| \leq 2^{-b}.$$

For case (b), analogously define, for input x, $w'(x,u) := c \cdot w(x,u)$ and $v'(x,u,b) := c \cdot v(x,u,b+a)$ where a is the smallest such that $c \leq 2^a$. Then, we have that

$$\begin{aligned} |w'(x,u) - v'(x,u,b)/2^b| &= \left| c \cdot w(x,u) - \frac{c \cdot v(x,u,b+a)}{2^{b+a}} \right| \\ &= \left| w(x,u) - \frac{v(x,u,b+a)}{2^{b+a}} + \cdots + w(x,u) - \frac{v(x,u,b+a)}{2^{b+a}} \right| \\ &\leq \left| c \cdot \frac{1}{2^{b+a}} \right| \leq \frac{1}{2^b}. \end{aligned}$$

We prove cases (c) and (d) for two functions f_1, f_2 on a common input x. This can be immediately generalized by a simple inductive argument. For (c), let f_1, f_2 the two functions associated with a weighted counting problem and let, for $i \in \{1,2\}$, $w_i(x,u)$ and $v_i(x,u,b)$ be the corresponding functions and their approximations in accordance with definition 6. For a fixed $u \in \{0,1\}^{p(|x|}$, let $w'(x,u) = w_1(x,u) + w_2(x,u)$ such that $f(x) = \sum_u w'(x,u)$. Define the approximation v' of w' to be $v'(x,u,b) = v_1(x,u,b+1) + v_2(x,u,b+1)$. Then, we have that

$$\begin{aligned} \left| w'(x,u) - \frac{v'(x,u,b)}{2^b} \right| &= \left| w_1(x,u) + w_2(x,u) - \frac{v_1(x,u,b+1) + v_2(x,u,b+1)}{2^{b+1}} \right| \\ &= \left| w_1(x,u) - \frac{v_1(x,u,b+1)}{2^{b+1}} + w_2(x,u) - \frac{v_1(x,u,b+1)}{2^{b+1}} \right| \\ &\leq \left| w_1(x,u) - \frac{v_1(x,u,b+1)}{2^{b+1}} \right| + \left| w_2(x,u) - \frac{v_1(x,u,b+1)}{2^{b+1}} \right| \\ &\leq \frac{1}{2^{b+1}} + \frac{1}{2^{b+1}} = \frac{1}{2^b}. \end{aligned}$$

For the product case of (d) define, for a fixed $u \in \{0,1\}^{p(|x|)}$, $w'(x,u) = w_1(x,u) \cdot w_2(x,u)$ such that $f(x) = \sum_u w_1(x,u) w_2(x,u)$. For $v'(x,u,b)$ that approximates w', define it as $v'(x,u,b) = v_1(x,u,b+a_2) \cdot v_2(x,u,b+a_1)$ where a_1 (respectively a_2) is the smallest number such that $w_1 \leq 2^{a_1}$ (respectively $w_2 \leq 2^{a_2}$) and let $a = \max\{a_1, a_2\}$ so that $w_1 + w_2 \leq 2^{a+1}$. We have that

$$\left| w'(x,u) - \frac{v'(x,u,b)}{2^b} \right| = \left| w_1(x,u) w_2(x,u) - \frac{v_1(x,u,b+a)}{2^{b+a+1}} \frac{v_2(x,u,b+a+1)}{2^{b+a+1}} \right|.$$

Using the property that $v_i(x,u,b)/2^b$ approximates $w_i(x,u)$ within an additive error of 2^{-b} (i.e., $w_i(x,u) - \frac{1}{2^b} \leq \frac{v_i(x,u,b)}{2^b} \leq w_i(x,u) + \frac{1}{2^b}$), we see that the above is less or equal to

$$\left| w_1(x,u) w_2(x,u) - \left(w_1(x,u) w_2(x,u) - \frac{(w_1(x,u) + w_2(x,u))}{2^{b+a+1}} + (\frac{1}{2^{b+a_1}})^2 \right) \right| \leq \frac{1}{2^b}.$$

For the concatenation cases, the proof uses in an immediate way the proofs for (c) and (d). Indeed, for $y \in \{0,1\}^{p(|x|)}$, let $w(\langle x,y \rangle, u) = w(x,u)$ and then it is enough to do the same for the approximation function v. Hence, we are back to cases (c) and (d) and the claim follows. □

Combining these results we can show that any multivariate polynomial in polynomially many variables (given as uniform weighted counting problems) with a degree bounded by a polynomial and with coefficients which are uniformly computed by another weighted counting problem is itself a weighted counting problem. The following theorem formalizes this statement.

Theorem 7. *Let f and c be functions computed by weighted counting problems and let q and r be polynomials. Define the function p for a given input $x \in \{0,1\}^n$ as*

$$p(x) = \sum_{(e_1,\ldots,e_{q(n)}) \in \{0,1,\ldots,r(n)\}^{q(n)}} c(\langle x, e_1, \ldots, e_{q(n)}\rangle) f(\langle x, 1\rangle)^{e_1} \cdots f(\langle x, q(n)\rangle)^{e_{q(n)}}.$$

Then p can be computed by a weighted counting problem.

Proof. This result follows from an application of Proposition 1. □

We make use of this result in the next section, which deals with the approximation of weighted counting problems by conventional counting problems.

2.5 Approximation

In the following we show that the approximation of general weighted counting problems up to a small additive approximation error is itself a #P-equivalent problem.

Theorem 8. *We are given a weighted counting problem in $\#P_{\mathbb{R}}$ defined by a polynomial p and a function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{R}$ as well as a number $b \in \mathbb{N}$. The task of approximating the resulting value of the given problem up to an additive approximation error of 2^{-b} is #P-equivalent.*

Proof. Hardness is trivial, since all problems in $\#\mathrm{P}_{\{0,1\}}$ fit the description. For a given $x \in \{0,1\}^\star$, we can solve the weighted counting problem given by the polynomial p and the integer weight function given by v where the last parameter is fixed to $p(|x|) + b$ with one call to a problem in $\#\mathrm{P}_{\mathbb{Z}}$, which is #P-equivalent due to Theorem 5, followed by a division by $2^{p(|x|)+b}$, so the reduction is affine. Let us call the function computed by this weighted counting problem f'. The approximation error of the resulting value with respect to the value obtained by the original weighted counting problem satisfies

$$\begin{aligned} |f(x) - f'(x)| &= \left| \sum_{u \in \{0,1\}^{p(|x|)}} w(x,u) - \sum_{u \in \{0,1\}^{p(|x|)}} v(x,u,p(|x|)+b)/2^{p(|x|)+b} \right| \\ &\leq \sum_{u \in \{0,1\}^{p(|x|)}} \left| w(x,u) - v(x,u,p(|x|)+b)/2^{p(|x|)+b} \right| \\ &\leq \sum_{u \in \{0,1\}^{p(|x|)}} 2^{-p(|x|)-b} \leq 2^{-b}, \end{aligned}$$

which concludes the proof. □

We can now use this result to show that weighted counting problems, for which the output (even if encoded as a fraction) is bounded polynomially in the input size, are in $\mathrm{FP}^{\#\mathrm{P}[1]}$.

Theorem 9. *We are given a weighted counting problem in $\#P_{\mathbb{R}}$ defined by a polynomial p and a function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{R}$. If the size of the output (even encoded as a fraction) is bounded by a polynomial $q(|x|)$, then the given weighted counting problem is in $FP^{\#P[1]}$.*

Proof. Using Theorem 8, we compute with a single call to a problem that is #P-equivalent a value $f'(x)$ with an additive approximation error of at most $2^{-2q(|x|)-2}$. We know that the actual value $f(x)$ is within the interval $\left[f'(x) - 2^{-2q(|x|)-2}, f'(x) + 2^{-2q(|x|)-2}\right]$ of size $2^{-2q(|x|)-1}$. In an interval of this size there is only one rational value which can be encoded by at most $q(|x|)$ many bits. Using the continued fraction expansions of the two interval endpoints we are able to retrieve this rational value efficiently [37], so the problem is in $\mathrm{FP}^{\#\mathrm{P}[1]}$ (a #P-equivalent class has the same power as #P if used as oracle). □

A similar result holds for the corresponding decision variant. In the following theorem we show that if the size of the output is bounded polynomially in the input size, the decision variant of the weighted counting problem is in PP.

Theorem 10. *We are given the decision variant of a weighted counting problem in $\#P_{\mathbb{R}}$ defined by a polynomial p, a function $w : \{0,1\}^\star \times \{0,1\}^\star \to \mathbb{R}$ and a threshold value $t \in \mathbb{N}$. If the size of the output of the corresponding weighted counting problem (even encoded as a fraction) is bounded by a polynomial $q(|x|)$, then the decision problem is in PP.*

Proof. This proof is also based on Theorem 8. As in the previous proof, we show that we can transform our problem to a problem with integer weights, such that the result of this problem divided by a certain integer gives us a good approximation for our original problem. Here we need two additional properties. First, we want to use an approximation from above, that means an approximation which is always at least as large as the exact value. In this way the approximate result is not smaller than the given threshold value if the exact result is not smaller than the given threshold. Second, we set the accuracy of the approximation such that the approximate value remains smaller than the given threshold value if the exact value is smaller than the threshold value. We can do this efficiently, since there is a certain gap between rational values whose denominators are bounded.

Take the function $v : \{0,1\}^\star \times \{0,1\}^\star \times \mathbb{N} \to \mathbb{Z}$ from Definition 6 in order to obtain a one-side approximation for the weights. For this purpose, simply compute one additional bit using the two-side approximation given by v and shift the result accordingly. In this way, the time bounds are not changed and yet we obtain a function $v' : \{0,1\}^\star \times \{0,1\}^\star \times \mathbb{N} \to \mathbb{Z}$ which approximates the weights and satisfies $0 \leq v'(x,u,b)/2^b - w(x,u) \leq 2^{-b}$ for all $x \in \{0,1\}^\star, u \in \{0,1\}^\star, b \in \mathbb{N}$.

For a given $c \in \mathbb{N}$, we can create the following problem in $\mathrm{PP}_{\mathbb{Z}}$. Take the same polynomial p as for the given problem. Additionally, take $v'(\cdot,\cdot,p(|x|)+c)$ as the integer weight function and $t' = t2^{p(|x|)+c}$ as the threshold value. Since the approximation error of the weights is always to the same side, it is easy to verify the following property: if the actual result of the original problem is at least as large as the threshold value t, then the result of the new problem is at least as large as the new threshold value t'.

It remains to show that if the actual resulting value of the original problem is smaller than t, then the resulting value of the new problem is also smaller than t'. This can be achieved by using an appropriate accuracy for the approximation, namely a value of $c = q(|x|) + 1$. The result of the new problem divided by $2^{p(|x|)+c}$ approximates the actual value with an one-side error of at most 2^{-c}. If the actual result of the original problem is smaller than t, then it differs from t by at least 2^{-c+1}, since the encoding of the denominator of the resulting value is bounded by $q(|x|)$. In that case the approximation of the actual result with an error of at most 2^{-c} is certainly smaller than t and therefore the result of the new problem is smaller than t'.

This concludes the proof, since we have shown that we can polynomially reduce the original problem to a single call of a problem in $\mathrm{PP}_{\mathbb{Z}} = \mathrm{PP}$. □

For problems from $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ to $\#\mathrm{P}_{\mathbb{R}}$, these approximate results are probably the best we can hope to obtain. The approximation results for weighted counting problems have a very interesting application regarding the multivariate polynomials which have been discussed in the last section. Of course, we can easily show that for a polynomially bounded output the exact evaluation of the multivariate polynomial is #P-equivalent and the corresponding decision variant is in PP. But instead of investigating such polynomials, it is much more interesting to focus on rational functions, that is, ratios of two such polynomials. The power of rational functions in the context of counting problems has been demonstrated in [7, 32], where the closure of PP under intersection and under truth-table reductions have been shown. While the (approximate) evaluation of rational functions seems to be more complex than solving a counting problem, the situation is different for the corresponding decision variant due to the observation that the sign of a rational function is equal to the sign of the product of numerator and denominator. If, in addition, the numerator and denominator are bounded from 0 by an exponentially small term, the sign can be decided in PP. The following theorem formalizes this statement.

Theorem 11. *Let p and q be two multivariate polynomials which can be computed by weighted counting problems, as used in Theorem 7. If there exists a polynomial t, such that for every input $x \in \{0,1\}^n$ we have $|p(x)| \geq 2^{-t(n)}$ and $|q(x)| \geq 2^{-t(n)}$, then the problem of deciding the sign of p/q is in PP.*

Proof. According to Theorem 7, p and q are both weighted counting problems. Due to Theorem 8, an additive approximation of p and q with error at most $2^{-t(n)-1}$ is a #P-equivalent problem. Moreover, such approximation preserves the signs of p and q. Therefore, the sign of p/q, which is the sign of pq, can be decided in PP, since deciding the sign of the approximated p and q are in PP (Theorem 10). □

3 The Closure of PP under Intersection and Truth-Table Reductions

In this section we give short and intuitive proofs of the closure of PP under intersection and truth-table reductions [7, 32]. These proofs show that it is more convenient to work in the more powerful class of weighted counting problems and performing an approximation at the very end to come back to the complexity of conventional counting problems. In this way, results about the approximation using rational functions can be directly used and additionally a lot of technicalities can be avoided. At this point we would like to emphasize that we are using the exactly same ideas as in the original proofs [7, 32] in the context of weighted counting. The contribution here is mainly the simplification of the original proofs.

The basis for these results (as well as for the original results) is Newman's Theorem [45]. This theorem states that a rational function with a degree of m can approximate the absolute value of a number in the interval $[-1,1]$ with an additive error of at most $3e^{-\sqrt{m}}$. This result can be trivially extended to functions in $[-c,c]$ (with $c > 1$), but the additive error increases to at most $c \cdot 3e^{-\sqrt{m}}$. Moreover, the proof of this result is constructive and the polynomials used for the rational function have a nice closed form expression. For further details we refer to the original publication [45].

Theorem 12 (Theorem 11 in [7])**.** *The intersection of finitely many languages in PP is itself in PP.*

Proof. Let $L_1, L_2, \ldots, L_k$ be languages in PP and let $L = \cap_{i=1}^{k} L_i$ denote the intersection of these languages. Then there exist corresponding GapP functions $f_1, f_2, \ldots, f_k$ such that for any input $x \in \{0,1\}^n$ we have $f_i(x) \geq 0 \iff x \in L_i$ and a polynomial p such that $|f_i(x)| \leq 2^{p(n)}$, with index $i \in \{1,2,\ldots,k\}$. Note that $f_i(x) < 0 \iff f_i(x) \leq -1$. Due to Newman's Theorem [45], there exist for each $i \in \{1,2,\ldots,k\}$ polynomials q_i and r_i of degree $4p(n)^2$ such that the value of $|f_i|$ is approximated by r_i/q_i to within $2^{p(n)}3e^{-2p(n)}$. Hence $|r_i(x)/q_i(x) - f_i(x)| \leq 2^{p(n)}3e^{-2p(n)} < 3e^{-p(n)}$ if $x \in L_i$ and $r_i(x)/q_i(x) - f_i(x) \geq 2 - 2^{p(n)}3e^{-2p(n)} > 1.5$ if $x \notin L_i$. Using these rational functions, we create the following function for deciding the intersection of the given languages.

$$f(x) = 1 - \sum_{i=1}^{k} \left(\frac{r_i(x)}{q_i(x)} - f_i(x) \right) .$$

Note that $f(x) \geq 0 \iff x \in L$. Manipulating $f(x)$, we have $f(x) = r(x)/q(x)$, where

$$r(x) = \sum_{i=1}^{k} (f_i(x)q_i(x) - r_i(x)) \prod_{j \neq i} q_j(x) + \prod_{j=1}^{k} q_j(x) \quad \text{and} \quad q(x) = \prod_{j=1}^{k} q_j(x) .$$

Now, a rather sloppy lower bound of $2^{-2kp(n)}$ on the absolute values of $r(x)$ and $q(x)$ allows us to apply Theorem 11 (the bound holds by construction, from [45]), which concludes the proof. □

Theorem 13 (Theorem 5.4 in [32]). *PP is closed under polynomial-time truth-table reductions.*

Proof. We are given a polynomial time algorithm computing the input for the PP oracle calls, a polynomial time predicate g representing the truth-table and a PP oracle represented by the GapP function f. For a given input $x \in \{0,1\}^n$ let $x_1, x_2, \ldots, x_{t(n)}$ be the polynomially many inputs of at most polynomial size for the oracle calls. Then there exists a polynomial p such that for any input $x \in \{0,1\}^n$ and any index $i \in \{1,2,\ldots,t(n)\}$ we have $|f(x_i)| \leq 2^{p(n)}$. Due to Newman's Theorem there exist for each $i \in \{1,2,\ldots,t(n)\}$ polynomials q_i and r_i of degree $4p(n)^2$ such that the value of $|f_i|$ is approximated by r_i/q_i to within $2^{p(n)}3e^{-2p(n)} < 3e^{-p(n)}$. As in the proof of Theorem 12, we can use the function $r_i/q_i - f_i$ to flatten positive cases to (near) zero while expanding negative cases away from zero. Additionally, the truth-table predicate can be written as a polynomial in the results from the oracle calls. This follows from arithmetization properties of deterministic Turing machines [5], which were also used in the original proof of this theorem [32]. Inserting the rational approximations into this polynomial results in a multivariate rational function. Due to the accuracy of the rational approximation it is clear that the sign of this function can be used to decide membership with respect to the truth-table reduction. Finally, using a similar technique as in the proof of Theorem 12, all terms can be brought to the same denominator and a rather sloppy lower bound of $2^{-2t(n)p(n)}$ on the absolute values of the numerator and denominator allows us to apply Theorem 11, which concludes the proof. □

4 Quantum Computing

In this section we investigate the relationship between weighted counting and quantum computation. The most common computational model for quantum computations is the quantum Turing machine. Such a Turing machine is similar to a probabilistic Turing machine [2]. Instead of (positive) transition probabilities, which preserve the ℓ_1 norm, quantum Turing machines use transition amplitudes, which are complex numbers that preserve the ℓ_2 norm. During the execution of a probabilistic Turing machine, the machine is at any point of time in exactly one state (according to the underlying probability distribution). The situation for quantum Turing machines is different. The machine is in a superposition of different states, whose probabilities are defined as the squares of their amplitudes. The exact nature of the state can only be determined through measurements. Such measurements are performed at the end of the computation to reveal the result, but also intermediate measurements are possible. Measuring certain bits destroys the superposition of states in the following sense. Only states which are compatible with the outcome of the measurement survive and the amplitudes change accordingly to preserve the ℓ_2 norm. The acceptance probability of a quantum Turing machine on a given input is the probability with which the final measurements reveal an accepting configuration.

We begin this section by showing that the acceptance probability of a quantum Turing machine can be represented using weighted counting. If there are no intermediate measurements, then the acceptance probability of the given quantum Turing machine is the sum of the acceptance probabilities for accepting configurations. The acceptance probability for an accepting configuration is the absolute square of its amplitude and its amplitude is the sum of the amplitudes of computational paths leading to this state. Finally, we can efficiently compute the amplitude of a given computational path of the quantum Turing machine to any precision we need. The key observation is that we can now express the overall acceptance probability by summing over all pairs of computational paths. We call two computational paths compatible if they arrive at the same accepting configuration. The weight for a pair of computational paths that are compatible is the real part of the product of the amplitude of the first computational path and the complex conjugate of the amplitude of the second computational path. Here we may consider only the real part of that product, since the imaginary parts cancel each other out over the whole sum. The weight for a pair of non-compatible computational paths is just 0. To allow intermediate measurements, we have to slightly modify the definition of compatible computational paths, but the underlying idea does not change. In the following we state this result in a more formal way.

Theorem 14. *For a given polynomial time quantum Turing machine, the task of computing the acceptance probability on any input is a weighted counting problem in $\#P_{\mathbb{R}}$.*

Proof. We first focus on polynomial time quantum Turing machines without intermediate measurements. After that we show how the result can be extended to the more general case.

Let M be a quantum Turing machine without intermediate measurements and whose computational time is bounded from above by a polynomial t. For a given input x, let us denote the set of possible final configurations of M by C and the set of possible computational paths by P. Note that the cardinality of both sets is at most exponential in the input length and that elements from these sets can be efficiently enumerated. The amplitude for a configuration $c \in C$ is denoted by a_c and the amplitude for a computational path $p \in P$ is denoted by a_p. We can then write the acceptance probability of M for the given input x in the following way (we use the notation $p_1 \sim p_2$ to indicate that two computational paths p_1 and p_2 arrive at the same accepting configuration).

$$\begin{aligned}
\Pr(M \text{ accepts } x) &= \sum_{c \in C} \left|a_c^2\right| = \sum_{c \in C} \left| \sum_{\substack{p \in P \\ \text{arriving at } c}} a_p \right|^2 \\
&= \sum_{c \in C} \sum_{\substack{p_1, p_2 \in P \\ \text{both arriving at } c}} a_{p_1} \overline{a_{p_2}} \\
&= \sum_{\substack{p_1, p_2 \in P \\ p_1 \sim p_2}} a_{p_1} \overline{a_{p_2}} = \sum_{\substack{p_1, p_2 \in P \\ p_1 \sim p_2}} \mathrm{Re}\left(a_{p_1} \overline{a_{p_2}}\right) , \qquad (1)
\end{aligned}$$

where $\mathrm{Re}(\cdot)$ is the real part of a number. To compute the acceptance probability, we have to sum over all pairs of computational paths, as shown in Expression (1). If two computational paths arrive at the same accepting configuration, we call them compatible. The weight for two compatible computational paths is the real part of the product of the amplitude of the first path with the complex conjugate of the amplitude of the second part. The weight for two non-compatible paths is just 0. In order to approximate the product of the amplitudes of two computational paths up to a specified precision, we just have to multiply the transition amplitudes between all the computational steps of the two paths with a certain precision. That means we can express the acceptance probability of M as a weighted counting problem. This derivation in is straightforward, but to our best knowledge it had not been used before for showing a direct relationship between quantum computation and weighted counting.

Now we still have to show that the result also holds if we allow intermediate measurements. Observe that this case is not necessary because, as it was shown in [10, 3], we can always postpone the intermediate measurements to the end of the computation. For the sake of completeness, we include the standard (and not very hard) arguments here.

For this purpose we have to adapt our definition of compatible computational paths. We now say that two computational paths are compatible if they arrive at the same accepting configuration and if they have the same results for intermediate measurements. We capture this extended definition of two compatible computational paths p_1 and p_2 by the same notation of $p_1 \sim p_2$. The general formula is then

$$\Pr(M \text{ accepts } x) = \sum_{\substack{p_1, p_2 \in P \\ p_1 \sim p_2}} \mathrm{Re}\left(a_{p_1} \overline{a_{p_2}}\right) .$$

To show that this formula is indeed correct, we postpone the intermediate measurements using the following standard technique. Here instead of an intermediate measurement, a controlled *not* (or simply CNOT) operation is performed on the bits that have to be measured and special ancilla bits. Then at the very end a measurement on the ancilla bits is performed. We call this

quantum Turing machine M'. M and M' have the same acceptance probabilities and in particular they accept the same language. Now our extended definition of compatible computational paths respects the following property: Two computational paths are compatible for M according to our extended definition if and only if the corresponding two computational paths are compatible for M' according the basic definition. This shows that such result also holds for the more general case in which intermediate measurements are allowed. □

We point out that there are no specific assumptions about the quantum Turing machines of the previous proof. In particular, we do not impose any restrictions on the transition amplitudes or on the measurements used thereby. Based on Theorem 14, we are now able to give extremely short and intuitive proofs of the well-known facts that BQP $\subseteq$ PP and (the stronger version) that BQP $\subseteq$ AWPP.

Definition 13. *A language $L \subseteq \{0,1\}^\star$ is in BQP if it can be decided by a polynomial-time quantum Turing machine with an error probability of at most $1/3$.*

Theorem 15. *BQP $\subseteq$ PP.*

Proof. For any given language $L \in$ BQP, there exists a bounded-error quantum Turing machine M which decides L. We can now apply Theorem 14 to M and obtain a weighted counting problem approximating the acceptance probabilities of M. The (approximate) decision version corresponding to this weighted counting problem is in PP and can be used to decide L due to the gap between the acceptance probabilities for words inside and outside the language. □

For the second proof we need an alternative definition of the complexity class AWPP. In [29] the following definition for AWPP is given.

Definition 14. *A language $L \subseteq \{0,1\}^\star$ is in AWPP if and only if there exists a polynomial p and a GapP function g such that, for all $x \in \{0,1\}^\star$,*

$$\begin{aligned} x \in L &\Rightarrow 1 - c \leq g(x)/2^p \leq 1, \\ x \notin L &\Rightarrow 0 \leq g(x)/2^p \leq c, \end{aligned}$$

where $p = p(|x|)$ and c is a constant smaller than $1/2$.

We now give an equivalent definition based on weighted counting.

Theorem 16. *A language $L \subseteq \{0,1\}^\star$ is in AWPP if and only if there exists a function g corresponding to a weighted counting problem such that, for all $x \in \{0,1\}^\star$,*

$$\begin{aligned} x \in L &\Rightarrow 1 - c \leq g(x) \leq 1, \\ x \notin L &\Rightarrow 0 \leq g(x) \leq c, \end{aligned}$$

where c is a constant smaller than $1/2$.

Proof. Due to Definition 14, it is clear that for a language $L \in$ AWPP there exists a function g corresponding to a weighted counting problem with the above properties. Now let us assume that the above properties are fulfilled for a function g corresponding to a weighted counting problem. We can approximate the weighted counting problem using a GapP function and a power of 2 as a scaling factor. For a sufficiently good approximation we are then able to fulfill the requirements of Definition 14 with some constant smaller than $1/2$. □

This alternative definition can now be used for a simplified proof of BQP $\subseteq$ AWPP.

Theorem 17. *BQP $\subseteq$ AWPP.*

Proof. For any given language $L \in$ BQP, there exists a bounded-error quantum Turing machine M which decides L. By applying Theorem 14 to M, we obtain a weighted counting problem approximating the acceptance probability of M. Since M has a bounded-error probability, we can show $L \in$ AWPP due to Theorem 16. □

Using the framework of weighted counting, it might be also possible to give simplified proofs for other results in the field of quantum computation. As an example, we give a shorter proof for a recent result [59] in the remaining part of this section. In [59], it is shown among other results that quantum Turing machines with bounded error can be simulated in a time- and space-efficient way using randomized algorithms with unbounded error that have access to random access memory. The most important version of this result has been stated in the following way.

Theorem 18 (Theorem 1.1 in [59])**.** *Every language solvable by a bounded-error quantum algorithm running in time $t \geq \log n$ and space $s \geq \log n$ with algebraic transition amplitudes is also solvable by an unbounded-error randomized algorithm with random access memory running in time $\mathcal{O}(t \log t)$ and space $\mathcal{O}(s + \log t)$, provided t and s are constructible by a deterministic algorithm with the latter time and space bounds.*

Proof. Let M be a given quantum Turing machine. We use again the formulation of the acceptance probability as a weighted counting problem:

$$\Pr(M \text{ accepts } x) \;=\; \sum_{\substack{p_1,\, p_2 \,\in\, P \\ p_1 \sim p_2}} \operatorname{Re}\left(a_{p_1} \overline{a_{p_2}}\right).$$

Based on this formulation, we can build a randomized algorithm with unbounded error: the problem of computing the acceptance probability of M is in $\#\mathrm{P}_{\mathbb{R}}$. The corresponding decision version is in $\mathrm{PP}_{\mathbb{R}}$ and a proper approximation resides in PP, which is a randomized algorithm with unbounded error. To show the desired time and space bounds, we have to bound the running time and the space used for the computation of each of the summands. For this purpose, one can use approximations of the constant number of transition amplitudes of M with a precision of $\mathcal{O}(\log t)$ bits. According to [9], the resulting quantum Turing machine is a good approximation to M. Note that the result was originally stated for quantum Turing machines without intermediate measurements, but with a similar argument as in the proof of Theorem 14, this result can be extended to the more general case. We now continue our investigations with this machine. We basically follow the proof of [59], but we employ our new terminology, which can help to state some parts in a much simpler way.

In a preprocessing step, we compute the (constant number of) transition amplitudes of M up to a precision of $\mathcal{O}(\log t)$ bits. This can be done for algebraic transition amplitudes in a computational time of $\mathcal{O}(\mathrm{polylog}(t))$ and space $\mathcal{O}(\log t)$. Due to the random access of the randomized machine, we can retrieve these values efficiently whenever they are required. Now the idea is to simulate all pairs of computational paths in parallel, step by step. Again due to the random access of the randomized machine, we can jump between the two computational paths without an additional overhead. For the moment we additionally keep track of the transition amplitudes of the paths. This causes some overhead that would invalidate the overall results, but we demonstrate how to avoid such overhead later. Whenever a quantum measurement occurs, we check if the two computational paths would result in the same measurement. If this is indeed the case, then we continue with the simulation, otherwise we stop and assign to this path a weight of 0. At the end, we check if both paths have arrived at the same accepting configuration. If that is the case, then we assign the real part of the product of the amplitude of the first path with the complex conjugate of the amplitude of the second part as the weight, otherwise we use a weight of 0. The resulting randomized algorithm shows the desired behavior, but does not respect the required time and space bounds due to the overhead needed to keep track of the transition amplitudes of the two computational paths. In the following we show how to fix this issue.

In order to avoid the overhead needed to keep track of the transition amplitudes of the two computational paths, one must realize that it is not even necessary to keep track of the weights throughout the whole computation. Instead, we may split the amplitudes in positive/negative and real/imaginary parts and perform random branchings at every computational step, instead of multiplying with the transition amplitudes. For that to work, the resulting weights of all the paths have to be adapted accordingly. If we look at time and space requirements of this approach, we

see that, at each computational step, we have to generate $\mathcal{O}(\log t)$ random bits for the branching, which results in a computational time of $\mathcal{O}(t \log t)$ and a space of $\mathcal{O}(s + \log t)$, as desired. □

We shall emphasize that these bounds rely heavily on the fact that random access to memory is granted to the randomized algorithm. Otherwise, we would obtain slightly weaker bounds for the general case. With a more elaborate approach it would still be possible to obtain the same bounds without random access to memory if the number of quantum measurements is $\mathcal{O}((\log t)^2)$. In this case, we have to move the pre-computed transition amplitudes during the simulation process to be always able to access them efficiently. This approach does not yield any overhead. Additionally, we simulate each of the computational paths until $\mathcal{O}(\log t)$ measurements occur. We must keep track of the measurement results and switch to the other path. In total, there are at most $\mathcal{O}(\log t)$ switches which cause a time overhead of $\mathcal{O}(s)$, which can be bounded from above by $\mathcal{O}(t)$.

5 Further Applications of Weighted Counting

In this section we take the previous discussions and results into two very relevant problems. First we discuss on inferences in the so called probabilistic graphical models, which appear in abundance in artificial intelligence, data mining and machine learning. Then we talk about stochastic combinatorial optimization problems, which represent a very important class of problems in operations research, with applications in numerous fields. Our approach focuses on using the complexity results presented so far to simplify or even to prove new complexity results for these problems.

5.1 Probabilistic Graphical Models

We present an application of the previously discussed complexity results to prove that *predictive inferences in probabilistic graphical models* (and especially in Bayesian networks [40, 47]) is #P-equivalent (predictive inferences are also called *belief updating* in this context). Historically, the community working with probabilistic graphical models has been used to cite papers that only partially resolve this question [42] [43], [51]. The most cited work is due to Roth [51], where hardness for #P is demonstrated, but membership is only superficially discussed and no formal proof is given. This issue is acknowledged many years later by Kwisthout [41], who in an attempt to close this question proves that predictive inference is in the so-called #P *modulo normalization* class. This situation is indeed inevitable, because the output here should be a probability value, so the problem cannot be in #P. However, using our terminology and results, we can state that such problems are #P-equivalent. By applying the results of this paper, we obtain a simple alternative proof to (and yet more powerful than) [41], as we discuss in the sequel of this section.

Definition 15 (Probabilistic Graphical Model (PGM))**.** *Let $\mathcal{J} = \{1, \ldots, n\}$ and $X_{\mathcal{J}} = (X_1, \ldots, X_n)$ be a vector of discrete random variables, $\mathcal{J}_1, \ldots, \mathcal{J}_m$ be a collection of index sets satisfying $\mathcal{J}_1 \cup \cdots \cup \mathcal{J}_m = \mathcal{J}$, and $\mathcal{P} = \{\phi_1, \ldots, \phi_m\}$ be a set of (explicitly represented, for instance in a table) functions over vectors $X_{\mathcal{J}_1}, \ldots, X_{\mathcal{J}_m}$ to non-negative rational numbers, respectively. We call $\mathcal{P}$ a* probabilistic graphical model *for $X_{\mathcal{J}}$ if the functions in $\mathcal{P}$ specify a joint probability distribution over assignments $x_{\mathcal{J}} \in X_{\mathcal{J}}$ by*

$$\Pr(X_{\mathcal{J}} = x_{\mathcal{J}}) = \frac{1}{Z} \prod_{i=1,\ldots,m} \phi_i(x_{\mathcal{J}_i}), \tag{2}$$

where $Z = \sum_{x_{\mathcal{J}} \in X_{\mathcal{J}}} \prod_{i=1,\ldots,m} \phi_i(x_{\mathcal{J}_i})$ is a normalizing value known as the partition function *and where the values $x_{\mathcal{J}_i}$ are compatible with $x_{\mathcal{J}}$. If the functions in $\mathcal{P}$ are not explicitly represented but can be computed in polynomial time in the size of the input, then we call $\mathcal{P}$ a* generalized probabilistic graphical model.

Note that the well-known Markov random fields as well as Bayesian networks [47] are nicely encompassed as subcases. For instance, a Bayesian network is a probabilistic graphical model that satisfies the following properties:

(i) it has exactly one ϕ_i for each X_i;

(ii) $\mathcal{J}_i$ must be such that $i \in \mathcal{J}_i$, and such that $j \notin \mathcal{J}_i$ whenever $j > i$;

(iii) $\sum_{x_i \in X_i} \phi_i(x_{\mathcal{J}_i}) = 1$.

Such restrictions naturally imply $Z = 1$ and induce conditional stochastic independence among variables $X_{\mathcal{J}}$ of the domain.

The *predictive inference* task can be succinctly defined as follows.

Definition 16 (Predictive inference in PGMs)**.** *Given a probabilistic graphical model $\mathcal{P}$ defined by functions $\{\phi_1, \ldots, \phi_m\}$ over discrete random variables $X_{\mathcal{J}} = (X_1, \ldots, X_n)$, predictive inference (also known as the computation of the partition function) is the task of computing $Z = \sum_{x_{\mathcal{J}} \in X_{\mathcal{J}}} \prod_{i=1,\ldots,m} \phi_i(x_{\mathcal{J}_i})$.*

Predictive inference turns out to be a general task that can be used to compute the probability value $\Pr(X_{\mathcal{J}'} = x_{\mathcal{J}'})$ for any event $x_{\mathcal{J}'}$ of interest: One has simply to take their specification of the Bayesian network and include into it the indicator functions $\prod_{j \in \mathcal{J}'} \psi_{x_j}(X_j = x_j)$. It is not hard to check that Z equals to $\Pr(X_{\mathcal{J}'} = x_{\mathcal{J}'})$ for a Bayesian network that is extended with these new indicator functions. Because of that, we call this probabilistic graphical model where Z equals to a probability value of interest as *queried Bayesian network*.

The complexity of predictive inference depends on how the input is encoded. When all functions in $\mathcal{P}$ are given by numbers directly encoded in the input, it is easy to show that the output has size that is polynomial in the input size (one could multiply all rational numbers in the input by their least common denominator in order to achieve an input defined solely by integers – this is done in polynomial time because all numbers are explicitly given). Hence the following theorem holds for any probabilistic graphical model, including queried Bayesian networks.

Theorem 19. *Given a probabilistic graphical model defined by $\mathcal{P}$, the predictive inference is #P-equivalent.*

Proof. Let D be the set of all denominators appearing in the images of any ϕ_i. Let $d = \prod_{r \in D} r$. Define new functions $\phi'_i = \phi_i \cdot d$. Now all input values are integer and so are all weights. This is clearly a problem in $\#\mathrm{P}_{\mathbb{N}}$. After the result is obtained, we divide the output by $d^{|X_{\mathcal{J}}|}$ and the result of the original problem is obtained. Hardness is obtained by reducing majority propositonal satisfiability (MAJSAT), as previously done in the literature, see for example [21]: create one X_i for each Boolean variable and a corresponding ϕ_i over it with image equals to $1/2$. Create a Boolean circuit representing the propositional formula using auxiliary Boolean variables X_j and functions ϕ_j, one for each time that a Boolean operator $\neg$, $\wedge$, $\vee$ appears in the propositional formula. Take the last variable X_n which subsumes the circuit output and attach to it the function ϕ_n which is an indicator of $X_n = \text{true}$. It is easy to see that Z is the proportion of satisfying assignments of the original propositional formula, as the circuit will evaluate to zero for every non-satisfying assignment and to $1/2^N$ for each satisfying assignment (with N the number of Boolean variables of MAJSAT). □

In the case of the generalized probabilistic graphical models, where functions are not given explicitly, we might end up with a large output size, but we can nevertheless show membership in $\#\mathrm{P}^{(\mathrm{pt})}_{\mathbb{Q}}$, which is straightforward.

Theorem 20. *Given a generalized probabilistic graphical model defined by $\mathcal{P}$, the predictive inference is $\#P^{(\mathrm{pt})}_{\mathbb{Q}}$-complete.*

Proof. Take the computation paths to correspond to values $x_{\mathcal{J}} \in X_{\mathcal{J}}$ and define the weight w of a path $x_{\mathcal{J}}$ to be the product $\prod_i \phi_i(x_{\mathcal{J}_i})$ of rational functions in $\mathcal{P}$. The sum of rational weights gives exactly the desired value of Z, so the problem is in $\#\mathrm{P}^{(\mathrm{pt})}_{\mathbb{Q}}$. Hardness comes from the fact that we can use a single ϕ_1 representing the weights of the $\#\mathrm{P}^{(\mathrm{pt})}_{\mathbb{Q}}$ problem. In this case, Z is exactly the sum of weights of the weighted counting problem with rational values. □

Regarding the decision version of *predictive inference*, where one queries whether Z is greater than a given rational threshold, the membership in PP is often attributed to Littman et al. [42], where membership of a similar (yet not equal) problem, namely probabilistic acyclic planning, is shown by the construction of a non-deterministic Turing machine with probability of acceptance greater than half. Such result, if manipulated properly, implies membership for the predictive inference in probabilistic graphical models too, but it is valid only in cases where the encoding of the instances satisfy some (restrictive) properties. This issue makes that result only partially satisfactory. Recently, Kwisthout [41] settles the membership of predictive inference in PP for queried Bayesian networks, but it does not extend to the generality of probabilistic graphical models as defined here. Hence, we obtain a stronger membership result, because we require only the output size to be polynomially bounded in the input size. Moreover, this is not restricted to queried Bayesian networks but works for any probabilistic graphical model, including those with functions that are parametric and shortly encoded (as long as they can be well approximated in polynomial time). In summary, results here lead to a simple proof that generalizes previous results for this problem [21, 41, 42, 43].

Theorem 21. *Given a generalized probabilistic graphical model defined by $\mathcal{P}$ such that its output size is known to be polynomial in the input size, the decision version of predictive inference is in PP.*

Proof. By applying Theorems 20 and 10, the result follows. □

Corollary 2. *Given a probabilistic graphical model defined by $\mathcal{P}$, the decision version of predictive inference is in PP.*

Proof. A probabilistic graphical model has explicitly represented functions, so its output size is known to be polynomial in the input size. Hence Theorem 21 suffices to achieve the desired result. □

Of greater interest is the *conditional predictive inference*.

Definition 17 (Conditional Predictive inference in PGMs)**.** *Given a probabilistic graphical model $\mathcal{P}$ defined by functions $\{\phi_1, \ldots, \phi_m\}$ over discrete random variables $X_{\mathcal{J}} = (X_1, \ldots, X_n)$, conditional predictive inference (also known as posterior probability computation) is the task of computing $\Pr(X_{\mathcal{J}'} = x_{\mathcal{J}'} \mid X_{\mathcal{J}''} = x_{\mathcal{J}''})$, for given instantiations $x_{\mathcal{J}'} \in X_{\mathcal{J}'}$ and $x_{\mathcal{J}''} \in X_{\mathcal{J}''}$.*

The computation of conditional predictive inference can be accomplished by two calls of predictive inference to obtain $\Pr(X_{\mathcal{J}'} = x_{\mathcal{J}'} \wedge X_{\mathcal{J}''} = x_{\mathcal{J}''})$ and $\Pr(X_{\mathcal{J}''} = x_{\mathcal{J}''})$, where probability values Pr are defined as in Expression (2). Even if the corresponding decision problem for conditional predictive inference is in PP, we are not aware of such proof (numerous papers cite sources which do not formally prove such result). Hence, a proof is given here.

Theorem 22. *Given a generalized probabilistic graphical model defined by $\mathcal{P}$ such that its output size is known to be polynomial in the input size, the decision version of conditional predictive inference is in PP.*

Proof. First of all, we can decide whether $\Pr(X_{\mathcal{J}''} = x_{\mathcal{J}''}) > 0$ (a call to PP suffices). If so, then let q be a rational given as the threshold. We want to decide whether

$$
\begin{aligned}
&\frac{\Pr(X_{\mathcal{J}' \cup \mathcal{J}''} = x_{\mathcal{J}' \cup \mathcal{J}''})}{\Pr(X_{\mathcal{J}''} = x_{\mathcal{J}''})} > q \iff \\
\iff &\Pr(X_{\mathcal{J}'} = x_{\mathcal{J}'} \wedge X_{\mathcal{J}''} = x_{\mathcal{J}''}) - q \cdot \Pr(X_{\mathcal{J}''} = x_{\mathcal{J}''}) > 0 \\
\iff &\sum_{y_{\mathcal{J}' \cup \mathcal{J}''} \in X_{\mathcal{J}' \cup \mathcal{J}''}} \Pr(X_{\mathcal{J}'} = y_{\mathcal{J}'} \wedge X_{\mathcal{J}''} = y_{\mathcal{J}''}) \left[\mathcal{I}(y_{\mathcal{J}' \cup \mathcal{J}''} = x_{\mathcal{J}' \cup \mathcal{J}''}) - q \cdot \mathcal{I}(y_{\mathcal{J}''} = x_{\mathcal{J}''})\right] > 0 \\
\iff &\sum_{y_{\mathcal{J}} \in X_{\mathcal{J}}} \Pr(X_{\mathcal{J}} = y_{\mathcal{J}}) \left[\mathcal{I}(y_{\mathcal{J}' \cup \mathcal{J}''} = x_{\mathcal{J}' \cup \mathcal{J}''}) - q \cdot \mathcal{I}(y_{\mathcal{J}''} = x_{\mathcal{J}''})\right] > 0 \\
\iff &\sum_{y_{\mathcal{J}} \in X_{\mathcal{J}}} w(y_{\mathcal{J}}) > 0\,,
\end{aligned}
$$

where $\mathcal{I}(\cdot)$ is the indicator function. Hence we can see the problem as an instance of $\mathrm{PP}_{\mathbb{Q}}^{(\mathrm{pt})}$ with weights $w(y_{\mathcal{J}}) = 0$ whenever $y_{\mathcal{J}}$ is not compatible with $x_{\mathcal{J}''}$, $w(y_{\mathcal{J}}) = -q \cdot \Pr(X_{\mathcal{J}} = y_{\mathcal{J}})$ if $y_{\mathcal{J}}$ is compatible with $x_{\mathcal{J}''}$ but not with $x_{\mathcal{J}'}$, and finally $w(y_{\mathcal{J}}) = (1-q) \cdot \Pr(X_{\mathcal{J}} = y_{\mathcal{J}})$ if $y_{\mathcal{J}}$ is compatible with both $x_{\mathcal{J}'}$ and $x_{\mathcal{J}''}$. Under the assumptions of this theorem (about bounded output), we can use Theorem 10 to show that this decision is in PP. Hence we need at most two adaptive calls of PP, which equals to PP itself. □

Corollary 3. *Given a probabilistic graphical model defined by $\mathcal{P}$, the decision version of conditional predictive inference is in PP.*

Proof. The same argument as in Corollary 2 suffices. □

We have refrained from discussing PP-hardness for these problems, since such result is very well-established, see for instance the book of Darwiche on the topic [21] (the reduction presented in the proof of Theorem 19 above can be used).

5.2 Stochastic Combinatorial Optimization

In this section we present another application of the complexity results, this time in the context of discrete two-stage stochastic combinatorial optimization problems [53].

Two-stage stochastic combinatorial optimization problems contain uncertainty in terms of stochastic data in their problem formulation. This uncertainty can for example be given by probability distributions over events of the domain. We call a specific realization of the random events a scenario and we assume that the number of different scenarios is bounded exponentially in the input size. We further assume that these scenarios can be enumerated efficiently and that the probability that a given scenario occurs can be computed efficiently in the sense of Definition 6. In the first stage a decision is made solely based on the given information, without knowing the actual realizations of the random events. This first-stage decision imposes certain costs. In the second stage, after the realization of the random events is revealed, another decision has to be taken based on the first-stage decision and on the revealed information. This second-stage decision can for example be used to guarantee feasibility of the final solution or to react to certain random events. The second-stage decision causes additional costs which are usually called recourse costs. The overall goal is to find a solution for the given two-stage stochastic combinatorial optimization problem which minimizes the total expected costs, which is defined by the costs of the first-stage decision plus the expected costs of the second-stage decision.

The formal definition of the model is given below.

Definition 18 (2-Stage Optimization)**.** *We are given a probability distribution over all possible realizations of the data, called scenarios and denoted by $\mathcal{S}$, and we construct a solution in two stages:*

1. *First, we may take some decisions to construct an initial part of the solution, x, incurring a cost of $c(x)$.*
2. *Then some scenario $A \in \mathcal{S}$ is realized according to the distribution, and in the second-stage we may augment the initial decisions by taking recourse actions $y(A)$, (if necessary) incurring some cost $f_A(x, y(A))$.*

The goal is to choose the initial decisions so as to minimize the expected total cost

$$c(x) + E_{A \in \mathcal{S}}[f_A(x, y(A))] .$$

Dyer and Stougie [25] have shown that discrete two-stage stochastic combinatorial optimization problems are #P-hard in general. In order to obtain this result, they make use of some artificial stochastic combinatorial optimization problems. The result has then been strengthened in [22, 60, 61], where #P-hardness has been shown for a practically relevant stochastic vehicle routing problem. These results are both imposing lower bounds on the computational complexity of

stochastic combinatorial optimization problems. Here we complement them with upper bounds for the computational complexity of many tasks related to discrete two-stage stochastic combinatorial optimization problems.

In this context the actual solution for a two-stage stochastic combinatorial optimization problems is usually the first-stage decision. This will become more clear with the following additional assumptions. Assume now that the costs for a first-stage decision can be computed in polynomial time and that for a given first-stage decision and a given scenario the corresponding recourse costs can also be computed in polynomial time. Given a solution we can compute the expected costs in the following way: By definition we are able to compute the costs imposed by the first-stage decision in polynomial time. We then enumerate the at most exponentially many scenarios and add to the total costs for each scenario the recourse costs of this scenario multiplied with the probability that this scenario occurs. Using Theorem 8 we can prove the following result for the evaluation of solutions.

Theorem 23. *We are given a discrete two-stage stochastic combinatorial optimization problem (respecting our assumptions) and a value $b \in \mathbb{N}$. The task of computing the expected costs for a solution up to an additive error of 2^{-b} is #P-equivalent.*

For most of the discrete two-stage stochastic combinatorial optimization problems it holds that the expected costs for any solution can be encoded using at most polynomially many bits in the input size. In fact, we are not aware of any problem of practical relevance in which this is not the case. Using this additional assumption we can prove the following results.

Theorem 24. *We are given a discrete two-stage stochastic combinatorial optimization problem (respecting our extended assumptions). Then the following results regarding the computational complexity of different computational tasks related to the given problem hold:*

(a) The task of computing the expected costs for a solution is #P-equivalent.

(b) The problem of deciding if a given solution has expected costs of at most $t \in \mathbb{Q}$ is in PP.

(c) The problem of deciding if a solution with expected costs bounded by $t \in \mathbb{Q}$ exists is in $NP^{\#P[1]}$.

(d) The problem of computing a solution with minimum expected costs can be solved in polynomial time with access to an $NP^{\#P[1]}$ oracle.

These results open some interesting paths for further research. First of all, they describe upper bounds for the complexity of various computational tasks related to two-stage stochastic combinatorial optimization. In [61, 60], it has already been shown that the upper bound of the first result in Theorem 24 is tight for a practically relevant problem. It remains to answer whether there are also practically relevant problems whose decision/optimization variants match the corresponding upper bounds given here.

6 Conclusions

We have presented a structured view on weighted counting. We have shown that weighted counting problems are a natural generalization of counting problems and that in many cases the computational complexity of weighted counting problems corresponds to that of conventional counting problems. The computational complexity of decision problems in $\mathrm{PP}_{\mathbb{Q}}^{(\mathrm{pt})}$, where the size of the output of the associated $\#\mathrm{P}_{\mathbb{Q}}^{(\mathrm{pt})}$ problem is not necessarily polynomially bounded in the input size, remains an interesting open problem. As for conventional counting problems, it is also of great interest to improve our understanding of the (polynomial time) approximability and inapproximability of weighted counting problems.

Additionally, we have seen that weighted counting problems arise in many different fields. Using the framework of weighted counting we could give more intuitive and simpler proofs for known

results regarding counting problems and quantum computation. We could even obtain new results regarding probabilistic graphical models and two-stage stochastic combinatorial optimization based on weighted counting. Finally, we believe that our results regarding weighted counting have many more applications in other situations and fields and we hope that our structured approach to weighted counting might help in revealing such relations in the near future.

7 Acknowledgments

The second author's research has been supported by the *Swiss National Science Foundation* as part of the *Early Postdoc.Mobility* grant P1TIP2_152282 and the third author's research has been supported by the *Swiss National Science Foundation* as part of the *Early Postdoc.Mobility* grant P2TIP2_152293.

References

[1] S. Aaronson, Quantum computing, postselection, and probabilistic polynomial-time, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 461 (2063) (2005) 3473–3482. doi:10.1098/rspa.2005.1546.

[2] S. Aaronson, Quantum Computing since Democritus, Cambridge University Press, 2013.

[3] L. Adleman, J. DeMarrais, M. Huang, Quantum computability, SIAM J. Comput. 26 (5) (1997) 1524–1540.

[4] S. Arora, B. Barak, Computational Complexity - A Modern Approach, Cambridge University Press, 2009.

[5] L. Babai, L. Fortnow, Arithmetization: A new method in structural complexity theory, Computational Complexity 1 (1) (1991) 41–66.

[6] R. Beigel, Closure properties of GapP and #P, in: ISTCS, 1997, pp. 144–146.

[7] R. Beigel, N. Reingold, D. Spielman, PP is closed under intersection, J. Comput. Syst. Sci. 50 (2) (1995) 191–202.

[8] A. Ben-Dor, S. Halevi, Zero-one permanent is $\#p$-complete, a simpler proof, in: ISTCS, 1993, pp. 108–117.

[9] E. Bernstein, U. Vazirani, Quantum complexity theory, in: Proceedings of the twenty-fifth annual ACM STOC, ACM, 1993, pp. 11–20.

[10] E. Bernstein, U. V. Vazirani, Quantum complexity theory, SIAM J. Comput. 26 (5) (1997) 1411–1473.

[11] M. Bläser, R. Curticapean, Weighted counting of k-matchings is #w[1]-hard, in: D. Thilikos, G. Woeginger (Eds.), IPEC, Vol. 7535 of Lecture Notes in Computer Science, Springer, 2012, pp. 171–181.

[12] A. Bulatov, M. Dyer, L. Goldberg, M. Jalsenius, M. Jerrum, D. Richerby, The complexity of weighted and unweighted #csp, J. Comput. Syst. Sci. 78 (2) (2012) 681–688.

[13] A. A. Bulatov, V. Dalmau, Towards a dichotomy theorem for the counting constraint satisfaction problem, in: FOCS, IEEE Computer Society, 2003, pp. 562–571.

[14] J. Cai, P. Lu, M. Xia, Computational complexity of holant problems, SIAM J. Comput. 40 (4) (2011) 1101–1132. doi:10.1137/100814585.
URL http://dx.doi.org/10.1137/100814585

[15] J. Cai, P. Lu, M. Xia, The complexity of complex weighted boolean #csp, J. Comput. Syst. Sci. 80 (1) (2014) 217–236.

[16] J. Cai, X. Chen, Complexity of counting csp with complex weights, in: H. Karloff, T. Pitassi (Eds.), STOC, ACM, 2012, pp. 909–920.

[17] J. Cai, X. Chen, P. Lu, Non-negatively weighted #csp: An effective complexity dichotomy, in: IEEE Conference on Computational Complexity, 2011, pp. 45–54.

[18] F. G. Cozman, D. D. Maua, Probabilistic graphical models specified by probabilistic logic programs: Semantics and complexity results, JMLR Workshop and Conference Proceedings 52 (2016) 110–122, proceedings of the Eighth International Conference on Probabilistic Graphical Models.

[19] F. G. Cozman, D. D. Maua, The structure and complexity of credal semantics, in: Proceedings of the 3rd International Workshop on Probabilistic Logic Programming, 2016, pp. 3–14.

[20] F. G. Cozman, D. D. Maua, The Complexity of Bayesian Networks Specified by Propositional and Relational Languages, ArXiv e-printsarXiv:1612.01120.

[21] A. Darwiche, Modeling and Reasoning with Bayesian Networks, Cambridge University Press, 2009.

[22] C. P. de Campos, G. Stamoulis, D. Weyland, The computational complexity of stochastic optimization, in: Proceedings of International Symposium on Combinatorial Optimization (ISCO), 2014, pp. 173–185.

[23] A. Durand, M. Hermann, P. G. Kolaitis, Subtractive reductions and complete problems for counting complexity classes, Theor. Comput. Sci. 340 (3) (2005) 496–513.

[24] M. Dyer, L. Goldberg, M. Jerrum, The complexity of weighted boolean csp, SIAM J. Comput. 38 (5) (2009) 1970–1986.

[25] M. Dyer, L. Stougie, Computational complexity of stochastic programming problems, Math. Prog. 106 (3) (2006) 423–432.

[26] M. Dyer, L. A. Goldberg, C. Greenhill, M. Jerrum, The relative complexity of approximate counting problems, Algorithmica 38 (3) (2004) 471–500. doi:10.1007/s00453-003-1073-y.
URL http://dx.doi.org/10.1007/s00453-003-1073-y

[27] P. Faliszewski, L. Hemaspaandra, The complexity of power-index comparison, Theoretical Computer Science 410 (1) (2009) 101–107.

[28] T. Feder, M. Vardi, The computational structure of monotone monadic snp and constraint satisfaction: A study through datalog and group theory, SIAM J. Comput. 28 (1) (1998) 57–104.

[29] S. Fenner, PP-lowness and a simple definition of AWPP, Theory of Computing Systems 36 (2) (2003) 199–212.

[30] S. Fenner, L. Fortnow, S. Kurtz, Gap-definable counting classes, J. Comput. Syst. Sci. 48 (1) (1994) 116–148.

[31] S. Fenner, L. Fortnow, L. Li, Gap-definability as a closure property, Inf. Comput. 130 (1) (1996) 1–17.

[32] L. Fortnow, N. Reingold, PP is closed under truth-table reductions, Inf. Comput. 124 (1) (1996) 1 – 6.

[33] L. Fortnow, J. D. Rogers, Complexity limitations on quantum computation, J. Comput. Syst. Sci. 59 (2) (1999) 240–252.

[34] J. Gill, Computational complexity of probabilistic turing machines, SIAM J. Comput. 6 (4) (1977) 675–695.

[35] L. A. Goldberg, M. Jerrum, Inapproximability of the tutte polynomial, Inf. Comput. 206 (7) (2008) 908–929.

[36] F. Green, J. Kobler, K. Regan, T. Schwentick, J. Toran, The power of the middle bit of a #P function, J. Comput. Syst. Sci. 50 (3) (1995) 456 – 467.

[37] G. Hardy, E. Wright, An introduction to the theory of numbers, Oxford University Press, 1979.

[38] L. A. Hemaspaandra, S. Kosub, K. W. Wagner, The complexity of computing the size of an interval, in: F. Orejas, P. G. Spirakis, J. van Leeuwen (Eds.), ICALP, Vol. 2076 of Lecture Notes in Computer Science, Springer, 2001, pp. 1040–1051.

[39] F. Jaeger, D. L. Vertigan, D. J. A. Welsh, On the computational complexity of the Jones and Tutte polynomials, Mathematical Proceedings of the Cambridge Philosophical Society 108 (1) (1990) 35–53.

[40] D. Koller, N. Friedman, Probabilistic Graphical Models, MIT press, 2009.

[41] J. Kwisthout, The computational complexity of probabilistic inference, Tech. rep., Faculty of Science, Radboud University Nijmegen (2011).

[42] M. Littman, J. Goldsmith, M. Mundhenk, The computational complexity of probabilistic planning, J. Artif. Intell. Res. (JAIR) 9 (1998) 1–36.

[43] M. Littman, S. M. Majercik, T. Pitassi, Stochastic boolean satisfiability, J. Autom. Reasonin 27(3) (2001) 251–296.

[44] M. Minsky, Computation, Prentice-Hall, 1967.

[45] D. Newman, Rational approximation to $|x|$, Michigan Math. J. 11 (1) (1964) 11–14.

[46] A. Pagourtzis, S. Zachos, The complexity of counting functions with easy decision version, in: R. Kralovic, P. Urzyczyn (Eds.), MFCS, Vol. 4162 of Lecture Notes in Computer Science, Springer, 2006, pp. 741–752.

[47] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, San Mateo, 1988.

[48] A. Rafiey, J. Kinne, T. Feder, Dichotomy for Digraph Homomorphism Problems, ArXiv e-printsarXiv:1701.02409.

[49] K. Regan, T. Schwentick, On the power of one bit of a #P function, in: Proceedings of the Fourth Italian Conference on Theoretical Computer Science, World Scientific, 1992, pp. 317–329.

[50] J. Rosser, L. Schoenfeld, Approximate formulas for some functions of prime numbers, Illinois J. Math. 6 (1962) 64–94.

[51] D. Roth, On the hardness of approximate reasoning, Artif. Intell. 82 (1-2) (1996) 273–302.

[52] S. Saluja, K. V. Subrahmanyam, M. N. Thakur, Descriptive complexity of #p functions, J. Comput. Syst. Sci. 50 (3) (1995) 493–505.

[53] L. Stougie, M. Van Der Vlerk, Stochastic integer programming, Institute of Actuarial Sciences & Econometrics, University of Amsterdam, 1996.

[54] S. Toda, PP is as hard as the polynomial-time hierarchy, SIAM J. Comput. 20 (5) (1991) 865–877.

[55] S. Toda, M. Ogiwara, Counting classes are at least as hard as the polynomial-time hierarchy, SIAM J. Comput. 21 (2) (1992) 316–328.

[56] S. Toda, O. Watanabe, Polynomial time 1-turing reductions from #PH to #PH, Theor. Comput. Sci. 100 (1) (1992) 205–221. doi:10.1016/0304-3975(92)90369-Q.
URL http://dx.doi.org/10.1016/0304-3975(92)90369-Q

[57] L. Valiant, The complexity of computing the permanent, Theor. Comput. Sci. 8 (1979) 189–201.

[58] L. Valiant, The complexity of enumeration and reliability problems, SIAM J. Comput. 8 (3) (1979) 410–421.

[59] D. van Melkebeek, T. Watson, Time-space efficient simulations of quantum computations, Theory of Computing 8 (1) (2012) 1–51.

[60] D. Weyland, On the computational complexity of the probabilistic traveling salesman problem with deadlines, Theor. Comput. Sci. 540 (2014) 156–168.

[61] D. Weyland, R. Montemanni, L. Gambardella, Hardness results for the probabilistic traveling salesman problem with deadlines, in: Proceedings of ISCO 2012 - The 2nd International Symposium on Combinatorial Optimization, 2012.

[62] T. Yamakami, Analysis of quantum functions, Int. J. Found. Comput. Sci. 14 (5) (2003) 815–852.

[63] T. Yamakami, Approximate counting for complex-weighted boolean constraint satisfaction problems, Inf. Comput. 219 (2012) 17–38. doi:10.1016/j.ic.2012.08.002.

[64] T. Yamakami, A dichotomy theorem for the approximate counting of complex-weighted bounded-degree boolean csps, Theor. Comput. Sci. 447 (2012) 120–135.